

RATIONAL SYMPLECTIC FIELD THEORY FOR LEGENDRIAN KNOTS

LENHARD NG

ABSTRACT. We construct a combinatorial invariant of Legendrian knots in standard contact three-space. This invariant, which encodes rational relative Symplectic Field Theory and extends contact homology, counts holomorphic disks with an arbitrary number of positive punctures. The construction uses ideas from string topology.

1. INTRODUCTION

The theory of Legendrian knots plays a key role in contact and symplectic topology and has recently shown surprising connections to low dimensional topology; see [Etn05] for a survey of the subject. A key breakthrough in the study of Legendrian knots, and symplectic topology generally, was the introduction of Gromov-type holomorphic-curve techniques in the 1990s. This led in particular to the development of Legendrian contact homology, outlined by Eliashberg and Hofer [Eli98] and fleshed out famously by Chekanov [Che02] for standard contact \mathbb{R}^3 and later by others in more general setups (e.g., [EES05a, EES07, NT04, Sab03]). Besides applications to contact topology, Legendrian contact homology has been closely linked to standard knot theory (e.g., [Ng08]).

Contact homology is part of a much larger construction, Symplectic Field Theory (SFT), which was introduced by Eliashberg, Givental, and Hofer about a decade ago [EGH00]. The relevant portion of the SFT package for our purposes is a filtered theory for contact manifolds whose first order comprises contact homology. Somewhat more precisely, while contact homology counts holomorphic disks in the symplectization of a contact manifold with exactly one positive boundary puncture, SFT counts holomorphic curves with arbitrarily many positive punctures.

In the “closed” case (in the absence of a Legendrian or Lagrangian boundary condition), SFT is now fairly well understood, both algebraically and analytically, and has produced a number of spectacular applications in symplectic topology; see, e.g., [Eli07] and references therein. However, in the “relative” case that is the focus of this paper, much less is currently understood. In particular, technical problems involving bubbling of holomorphic curves have thus far prevented a formulation of SFT with Legendrian boundary condition even for the basic case of standard contact \mathbb{R}^3 . The development of contact homology for Legendrian knots involves two steps,

a fairly easy proof that $d^2 = 0$ and a more difficult invariance proof; it has proven surprisingly difficult to extend this to a reasonable algebraic setup for Legendrian SFT that even satisfies $d^2 = 0$, not to mention invariance.

In this paper, we will give an algebraic formulation, à la Chekanov [Che02], of Legendrian SFT for standard contact \mathbb{R}^3 ; this allows us to skirt the analytical issues that would usually beset the proofs of $d^2 = 0$ and invariance. We note that we present not the full Legendrian SFT, which would consider holomorphic curves of arbitrary genus and possibly marked points and gravitational descendants, but “rational” SFT, which counts only holomorphic disks.¹

The technique that we use to overcome the bubbling problems comes from string topology [CS]. Cieliebak and Latschev [CL], motivated by work of Fukaya, have developed a program for using string topology to deal with compactification issues in Legendrian SFT; see especially the appendix to [CL] jointly written with Mohnke. The program currently has significant unresolved technical issues, but one can avoid these issues in the case of \mathbb{R}^3 by using the combinatorial approach we employ here. On a related note, we remark without proof that a separate approach to Legendrian SFT, based on the cluster homology of Cornea and Lalonde [CL06], seems in the \mathbb{R}^3 case to yield the same theory as ours, or at least the commutative quotient that we call $(\hat{\mathcal{A}}^{\text{comm}}, d)$.

We now outline the mathematical content of this paper. In Section 2, we associate to any Legendrian knot in standard contact \mathbb{R}^3 a filtered version of a familiar structure from algebra, a curved dg-algebra, which itself is a type of a curved A_∞ algebra. Our particular filtered curved dg-algebra, which we call the *LSFT algebra* $(\hat{\mathcal{A}}, d)$ of the Legendrian knot, takes the following form: $\hat{\mathcal{A}}$ is the tensor algebra over \mathbb{Z} freely generated by two generators p, q for each Reeb chord, along with one more generator t and its inverse t^{-1} essentially encoding the homology of the knot. The map d on $\hat{\mathcal{A}}$ is a derivation

$$d = d_{\text{SFT}} + d_{\text{str}},$$

where d_{SFT} is an “SFT differential” obtained by counting rational holomorphic curves in the symplectization $\mathbb{R} \times \mathbb{R}^3$ with boundary on the Lagrangian cylinder over the knot and boundary punctures approaching Reeb chords at $\pm\infty$ in the distinguished \mathbb{R} direction, and d_{str} is a “string differential” encoding a string cobracket operation that glues trivial holomorphic strips to broken closed strings on the knot. The Hamiltonian that produces the SFT differential lives naturally in the quotient $\hat{\mathcal{A}}^{\text{cyc}}$ of $\hat{\mathcal{A}}$ by cyclic permutations but acts on $\hat{\mathcal{A}}$ as well.

The string differential is a necessary correction that accounts for the aforementioned bubbling and ensures a result analogous to $d^2 = 0$. More precisely, we have the following two main results.

¹This is a slight misuse of the term “rational” since we do not count genus-0 surfaces with more than one boundary component.

Theorem 1.1 (see Theorem 2.25). *The algebra $(\hat{\mathcal{A}}, d)$ associated to a Legendrian knot is a curved dg-algebra; that is, there is an element F_d of $\hat{\mathcal{A}}$ such that $d^2(x) = F_dx - xF_d$ for all $x \in \hat{\mathcal{A}}$.*

Theorem 1.2 (see Theorem 2.28). *$(\hat{\mathcal{A}}, d)$ is invariant under restricted Legendrian isotopies.*

Here “restricted” is a minor technical condition (see Definition 2.26) that we conjecture can be removed, but that in any case can still be used to produce an invariant of Legendrian knots under arbitrary Legendrian isotopies; see Corollary 2.29. It is possible that we can remove the “restricted” condition if we allow arbitrary equivalences of curved A_∞ algebras rather than the specific equivalences of LSFT algebras defined in Section 2.2, but we do not pursue this point in this paper.

The LSFT algebra has a filtration whose associated graded object, in the bottom filtration level, is Legendrian contact homology (cf. Remark 2.31). Theorems 1.1 and 1.2 contain Chekanov’s $d^2 = 0$ and invariance results for contact homology (Corollary 2.30).

One possible and desirable application of Legendrian SFT would be the construction of invariants of Legendrian knots that do not vanish for stabilized knots, which in some sense comprise “most” Legendrian knots. This could produce invariants of topological knots (which can be viewed as Legendrian knots modulo stabilization) and transverse knots (Legendrian knots modulo one particular stabilization), among other things. Contact homology famously vanishes under stabilization [Che02], but it was hoped for some time that Legendrian SFT would not. Unfortunately, rational Legendrian SFT, as constructed in this paper, also loses all interesting information under stabilization; see Appendix B. There is some hope that one could apply rational SFT to the double of a Legendrian knot [NT04] to obtain an interesting invariant, but this is unclear as yet. We note that the contact homology of the double of a stabilized knot contains no information [Ng01], but rational SFT may encode significantly more information.

We remark that we develop the theory over \mathbb{Z} , and a fair amount of work throughout the paper is devoted to keeping track of signs. In particular, we include an appendix that computes all possible sign rules, in some suitable sense, and shows that they are all equivalent. However, the entire theory works over $\mathbb{Z}/2$ as well as \mathbb{Z} , with the notable exception of invariance for cyclic and commutative complexes (Proposition 2.33), and the reader may find it easier to ignore all signs.

In this paper, we omit discussion of the relation between our algebraic version of rational Legendrian SFT and the more general, more geometric string-topology version, though we may return to this topic in the future. We also postpone concrete applications of the Legendrian SFT formalism presented here, such as the construction of an L_∞ structure on cyclic Legendrian contact homology, to a future paper.

The main results of this paper are contained in Section 2. Their proofs, some of which involve a discussion of a rudimentary version of string topology, occupy Sections 3 (for Theorem 1.1) and 4 (for Theorem 1.2). Appendices A and B deal with sign choices and triviality for stabilized knots, respectively.

Acknowledgments. I would like to give significant thanks to Yasha Eliashberg, whose many conversations with me about candidates for Legendrian SFT played a key role in the present work. In addition, the crucial catalyst for this paper was the September 2007 workshop “Towards Relative Symplectic Field Theory” sponsored by the American Institute of Mathematics, the NSF, the CUNY Graduate Center, and the Stanford Mathematical Research Center. I am deeply indebted to all of the workshop’s participants, particularly Mohammed Abouzaid, Frédéric Bourgeois, Kai Cieliebak, Tobias Ekholm, John Etnyre, Eleny Ionel, Janko Latschev, and Josh Sabloff, for their ideas and suggestions, and to Mikhail Khovanov for a separate illuminating conversation. The combinatorial version of Legendrian SFT presented here was largely formulated in discussions at the AIM workshop. I also thank the referee for helpful corrections and suggestions. This work is partially supported by the following NSF grants: DMS-0706777, FRG-0244663, and CAREER grant DMS-0846346.

2. THE SFT INVARIANT

In this section, we describe the algebraic object to be associated to a Legendrian knot, the LSFT algebra, and state the main “ $d^2 = 0$ ” and invariance results, though their proofs are deferred to Sections 3 and 4. The LSFT algebra is a special case of a familiar construction from homological algebra, the curved dg-algebra, whose salient features we review in Section 2.1. We then present the definition of an LSFT algebra in Section 2.2, followed by a combinatorial definition for the LSFT algebra associated to the xy projection of a Legendrian knot in Section 2.3. In Section 2.4, we discuss two quotient invariants derived from the LSFT algebra, the cyclic and commutative complexes.

2.1. Algebraic setup: curved dg-algebras. Throughout this section and the paper, we use the convention that the commutator on a graded associative algebra is $[x, y] = xy - (-1)^{|x||y|}yx$.

Definition 2.1. A *curved dg-algebra* consists of a triple (\mathcal{A}, d, F) , where:

- \mathcal{A} is a graded associative algebra over \mathbb{Z} ;
- $d : \mathcal{A} \rightarrow \mathcal{A}$ is a derivation, i.e., $d(xy) = (dx)y + (-1)^{|x|}x(dy)$, and d lowers degree by 1;
- F is a degree -2 element of \mathcal{A} (the *curvature*) for which $dF = 0$;

- for all $x \in \mathcal{A}$, $d^2(x) = [F, x]$.

A *filtered curved dg-algebra* is a curved dg-algebra with a descending filtration of subalgebras

$$\mathcal{A} = \mathcal{F}^0\mathcal{A} \supset \mathcal{F}^1\mathcal{A} \supset \mathcal{F}^2\mathcal{A} \supset \dots$$

with respect to which d is a filtered derivation and $F \in \mathcal{F}^1\mathcal{A}$.

Remark 2.2. Curved dg-algebras have been studied extensively in the literature, though sometimes under other names, e.g., CDG-algebra [Pos93] and Q -algebra [Sch99] (note however that the standard definition involves an algebra over a field rather than over \mathbb{Z}). In particular, a curved dg-algebra is essentially a special case of a curved (or “weak”) A_∞ algebra. A curved A_∞ algebra is a graded vector space V with multilinear maps $m_n : V^{\otimes n} \rightarrow V$ of degree $n - 2$ for all $n \geq 0$, satisfying the curved A_∞ relations

$$\sum_{i+j+k=n} (-1)^{i+j(1+i+|a_1|+\dots+|a_i|)} m_{i+k+1}(a_1, \dots, a_i, m_j(a_{i+1}, \dots, a_{i+j}), a_{i+j+1}, \dots, a_{i+j+k}) = 0$$

for $n \geq 0$. Except for the aforementioned discrepancy in base ring, a curved dg-algebra is a curved A_∞ algebra where $m_n = 0$ for $n \geq 3$; we then have $m_0 = F$, $m_1(a_1) = da_1$, and $m_2(a_1, a_2) = a_1a_2$, and the curved A_∞ relations become the relations in Definition 2.1, along with multiplicative associativity. For comparison, note that a usual A_∞ algebra is a curved A_∞ algebra with $m_0 = 0$, while a usual dg-algebra has $m_n = 0$ for all $n \neq 1, 2$.

A special case of morphisms of curved A_∞ algebras is the following.

Definition 2.3. A *morphism* of curved dg-algebras is a map $(\varphi, \alpha) : (\mathcal{A}, d, F) \rightarrow (\mathcal{A}', d', F')$, where:

- $\varphi : \mathcal{A} \rightarrow \mathcal{A}'$ is a graded algebra map;
- α is a degree -1 element of \mathcal{A}' ;
- $d'\varphi(\cdot) = \varphi d(\cdot) + [\alpha, \varphi(\cdot)]$;
- $F' = \varphi(F) + d\alpha + \alpha^2$.

A *filtered morphism* of filtered curved dg-algebras is a morphism for which φ respects the filtration and $\alpha \in \mathcal{F}^1\mathcal{A}'$.

It is easy to check that a composition of morphisms is a morphism, where we define $(\varphi', \alpha') \circ (\varphi, \alpha) = (\varphi' \circ \varphi, \alpha' + \varphi'(\alpha))$. There is an identity morphism $(\text{Id}, 0)$, and if (φ, α) is a morphism for which φ is an isomorphism, then $(\varphi^{-1}, -\varphi^{-1}\alpha)$ provides an inverse to (φ, α) .

We can now define chain homotopy and homotopy equivalence in the usual way. We state the definitions for filtered curved dg-algebras; there is an obvious analogue in the unfiltered case.

Definition 2.4. Two filtered morphisms of filtered curved dg-algebras $(\varphi, \alpha), (\varphi', \alpha') : (\mathcal{A}, d, F) \rightarrow (\mathcal{A}', d', F')$ are *chain homotopic* if $\alpha = \alpha'$ and there exists a filtered \mathbb{Z} -module map $H : \mathcal{A}' \rightarrow \mathcal{A}'$ of degree 1 such that

$$\varphi - \varphi' = Hd' + d'H.$$

A filtered morphism $(\varphi, \alpha) : (\mathcal{A}, d, F) \rightarrow (\mathcal{A}', d', F')$ is a *homotopy equivalence* if there exists a filtered morphism $(\varphi', \alpha') : (\mathcal{A}', d', F') \rightarrow (\mathcal{A}, d, F)$ such that $(\varphi', \alpha') \circ (\varphi, \alpha)$ and $(\varphi, \alpha) \circ (\varphi', \alpha')$ are each chain homotopic to the identity morphism.

We can now state a preliminary version of the main result of this paper; see Theorems 2.25 and 2.28 for the precise statements.

Theorem 2.5. *Rational SFT gives a map from Legendrian knots in \mathbb{R}^3 modulo Legendrian isotopy to filtered curved dg-algebras modulo homotopy equivalence.*

Because of the curvature term F , a curved dg-algebra does not typically comprise a complex. One can produce a complex and thus homology from a filtered curved dg-algebra in several ways. See Remark 2.7 for discussion of the associated graded complex, and Section 2.4 for the cyclic and commutative complexes.

2.2. Algebraic setup: LSFT algebras. The invariant we associate to a Legendrian knot is a particular type of filtered curved dg-algebra that we term an LSFT algebra. Besides being a specialization of the construction in the previous section, our definition of LSFT algebra generalizes (and contains) Chekanov's DGAs and stable tame isomorphisms from Legendrian contact homology.

Underlying an LSFT algebra is a (based) tensor algebra \mathcal{A} over \mathbb{Z} generated by $q_1, \dots, q_n, p_1, \dots, p_n, t, t^{-1}$; this is noncommutative and has sole relations $t \cdot t^{-1} = t^{-1} \cdot t = 1$. We consider $q_1, \dots, q_n, p_1, \dots, p_n$ to be distinguished generators that are included in the data of the LSFT algebra, where we view q_i and p_i as being paired together for $i = 1, \dots, n$, but the indices $1, \dots, n$ can be permuted without changing \mathcal{A} . Each generator of \mathcal{A} is \mathbb{Z} -graded with $|q_i| + |p_i| = -1$ for all i , and $|t| = -|t^{-1}| = -2r$ for some $r \in \mathbb{Z}$; this grading induces a grading on \mathcal{A} .

There is a filtration

$$\mathcal{A} = \mathcal{F}^0\mathcal{A} \supset \mathcal{F}^1\mathcal{A} \supset \mathcal{F}^2\mathcal{A} \supset \dots,$$

where $\mathcal{F}^k\mathcal{A}$ is generated by words containing at least k p 's. (Note that $\mathcal{F}^k\mathcal{A} = (\mathcal{F}^1\mathcal{A})^k$.) We will sometimes write $O(p^k)$ to denote an element of $\mathcal{F}^k\mathcal{A}$ (or $\mathcal{F}^k\hat{\mathcal{A}}$, defined below), and $x \equiv y \pmod{p^k}$ for $x = y + O(p^k)$.

Let $\hat{\mathcal{A}}$ be the “ p -adic completion” of \mathcal{A} consisting of possibly infinite sums $\sum_{k=0}^{\infty} z_k$ with $z_k \in \mathcal{F}^k\mathcal{A}$ for all k . That is, $\hat{\mathcal{A}}$ includes infinite sums in \mathcal{A} as long as for each k , all but finitely many terms in the sum do not lie in $\mathcal{F}^k\mathcal{A}$.

Then $\hat{\mathcal{A}}$ inherits from \mathcal{A} the structure of a graded algebra with filtration $\hat{\mathcal{A}} = \mathcal{F}^0\hat{\mathcal{A}} \supset \mathcal{F}^1\hat{\mathcal{A}} \supset \dots$.

Definition 2.6. An *LSFT algebra* is a filtered graded tensor algebra $\hat{\mathcal{A}} = \mathbb{Z}\langle q_1, \dots, q_n, p_1, \dots, p_n, t, t^{-1} \rangle$, as above, with a derivation² $d: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ satisfying the following conditions:

- (1) d has degree -1 and preserves the filtration;
- (2) $d(t) \in \mathcal{F}^1\hat{\mathcal{A}}$;
- (3) there is an element $F_d \in \mathcal{F}^1\hat{\mathcal{A}}$, the *curvature* of d , such that $d^2x = [F_d, x]$ for all $x \in \hat{\mathcal{A}}$.

We denote an LSFT algebra by (\mathcal{A}, d) , omitting the curvature F_d , which is uniquely determined by d .

Condition (3) ensures that $dF_d = 0$, since $[F_d, dx] = d^2(dx) = d(d^2x) = [dF_d, x] + [F_d, dx]$ for all $x \in \hat{\mathcal{A}}$; thus an LSFT algebra is a filtered curved dg-algebra in the sense of Section 2.1.

Remark 2.7 (The Chekanov–Eliashberg DGA). Given a curved dg-algebra $(\hat{\mathcal{A}}, d, F)$, one can consider the complex given by the associated graded object $\bigoplus_{i=0}^{\infty} \mathcal{F}^i\hat{\mathcal{A}}/\mathcal{F}^{i+1}\hat{\mathcal{A}}$ with the induced differential. In the case when $(\hat{\mathcal{A}}, d)$ is an LSFT algebra generated by $q_1, \dots, q_n, p_1, \dots, p_n, t, t^{-1}$, the $i = 0$ summand $(\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}, d)$ of the associated graded complex is generated by $q_1, \dots, q_n, t, t^{-1}$, with $d(t) = d(t^{-1}) = 0$.

This quotient $(\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}, d)$ is essentially Chekanov’s differential graded algebra (usually abbreviated DGA), also formulated by Eliashberg, that encodes Legendrian contact homology. Indeed, it will be clear from the definition of d in Section 2.3 that the differential on $\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}$, and in fact the entire associated graded object, counts precisely the same holomorphic disks as contact homology, namely disks with exactly one positive puncture. It should be noted, however, that $(\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}, d)$ is not precisely the same as the Chekanov DGA; see Remark 2.31 below.

Notation. We will sometimes want to treat the q ’s and p ’s together, and will use s to denote any q_j or p_j (or sometimes $t^{\pm 1}$ as well). The q ’s and p ’s are paired together, and we use $*$ to denote the pairing; that is, write $p_j^* = q_j$, $q_j^* = p_j$. We reserve w to mean a word in the q ’s, p ’s, and $t^{\pm 1}$.

If s is a q_j or p_j , then define $\{s, s^*\}$ to be $+1$ if s is a p and -1 if s is a q ; this is a special case of the SFT bracket to be defined in Section 3.1.

We next define a notion of equivalence between LSFT algebras, which is a special case of homotopy equivalence between filtered curved dg-algebras

²As in the previous section, a derivation is a \mathbb{Z} -linear map $d: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ such that $d(xy) = (dx)y + (-1)^{|x|}x(dy)$ for all $x, y \in \hat{\mathcal{A}}$ for which x is of pure degree. Note that, for an LSFT algebra, d necessarily satisfies $d(1) = 0$ and $d(t^{-1}) = -t^{-1} \cdot d(t) \cdot t^{-1}$.

(see Proposition 2.15 below). To do this, we introduce a specific family of curved dg-morphisms between LSFT algebras.

Definition 2.8. Let $\hat{\mathcal{A}}$ be an LSFT algebra (without its differential). An *elementary automorphism* of $\hat{\mathcal{A}}$ is a grading-preserving algebra automorphism ϕ of $\hat{\mathcal{A}}$ of one of the following forms:

- (1) $\phi(q_i) = \pm t^{\alpha_i} q_i t^{\beta_i}$, $\phi(p_i) = \pm t^{\gamma_i} p_i t^{\delta_i}$, $\phi(t) = t$, $\phi(t^{-1}) = t^{-1}$ for some integers $\alpha_i, \beta_i, \gamma_i, \delta_i$;
- (2) for some j , $\phi(q_i) = q_i$ for all $i \neq j$, $\phi(p_i) = p_i$ for all i , $\phi(t^{\pm 1}) = t^{\pm 1}$, and

$$\phi(q_j) = q_j + v + u$$

where $v \in \hat{\mathcal{A}}$ does not involve q_j and $u \in \mathcal{F}^1 \hat{\mathcal{A}}$;

- (3) for some j , $\phi(p_i) = p_i$ for all $i \neq j$, $\phi(q_i) = q_i$ for all i , $\phi(t^{\pm 1}) = t^{\pm 1}$, and

$$\phi(p_j) = p_j + v + u$$

where $v \in \hat{\mathcal{A}}$ does not involve p_j and $u \in \mathcal{F}^2 \hat{\mathcal{A}}$;

- (4) $\phi(q_i) = q_i$ and $\phi(p_i) = p_i$ for all j , and $\phi(t) = t + v$ for some $v \in \mathcal{F}^1 \hat{\mathcal{A}}$.

In the last three cases, we say that the elementary automorphism is *supported* at the generator of $\hat{\mathcal{A}}$ on which it is nontrivial: q_j for (2), p_j for (3), t for (4).

Implicit in the above definition is the following fact.

Lemma 2.9. *Each of the maps in Definition 2.8 is invertible.*

Proof. Maps of type (1) in the statement of Definition 2.8 are obviously invertible. Next consider a map ϕ of type (2). It suffices to show that ϕ is invertible if either $v = 0$ or $u = 0$, since in the general case, $\phi = \phi_1 \circ \phi_2$, where ϕ_1, ϕ_2 are supported on the same q_j and $\phi_1(q_j) = q_j + v$, $\phi_2(q_j) = q_j + \phi_1^{-1}(u)$ (note $\phi_1^{-1}(u) \in \mathcal{F}^1 \hat{\mathcal{A}}$). Now if $u = 0$, then ϕ is clearly invertible: $\phi^{-1}(q_j) = q_j - v$. If $v = 0$, define $\psi: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ by $\psi(x) = \phi(x) - x$ for all $x \in \hat{\mathcal{A}}$; then ψ increases filtration level by 1, and

$$\phi^{-1}(q_j) = q_j - u + \psi(u) - \psi \circ \psi(u) + \psi \circ \psi \circ \psi(u) - \dots$$

gives the inverse for ϕ .

The same proof works for a map of type (3).

Finally, suppose that ϕ is of type (4). We can define

$$\phi(t^{-1}) = t^{-1} - t^{-1} \cdot v \cdot t^{-1} + t^{-1} \cdot v \cdot t^{-1} \cdot v \cdot t^{-1} - \dots$$

and then ϕ is an algebra map on $\hat{\mathcal{A}}$; also, ϕ is invertible for the same reason as in case (2). \square

Definition 2.10. We say that LSFT algebras $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}, d')$ are related by a *basis change* if there is a sequence of elementary automorphisms of $\hat{\mathcal{A}}$ sending d to d' .

We remark that the composition ϕ of elementary automorphisms in Definition 2.10 yields a curved dg-morphism $(\phi, 0)$ between the curved dg-algebras given by $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}', d')$.

If $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}', d')$ are related by a basis change, then the quotient differential graded algebras $(\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}, d)$ and $(\mathcal{F}^0\hat{\mathcal{A}}'/\mathcal{F}^1\hat{\mathcal{A}}', d')$ are related by a tame isomorphism in the sense of Chekanov (see [Che02, ENS02] for the precise definition). Note that on the quotient level, any basis change fixes t and t^{-1} .

We need two more operations on LSFT algebras, gauge change and stabilization.

Definition 2.11. We say that LSFT algebras $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}', d')$ are related by a *gauge change* if there exists $\alpha \in \mathcal{F}^1\hat{\mathcal{A}}$ with $|\alpha| = -1$ such that

$$(1) \quad d'(x) = d(x) + [\alpha, x]$$

for all $x \in \hat{\mathcal{A}}$.

It is easy to check that if $(\hat{\mathcal{A}}, d)$ is an LSFT algebra, then (1) defines an LSFT algebra $(\hat{\mathcal{A}}', d')$ with $F_{d'} = F_d + d(z) + z^2$. Note that a gauge change is nothing more than a curved dg-morphism of the form (Id, α) .

Remark 2.12. Our notion of a gauge change coincides with the standard algebraic notion of changing by an inner derivation. One can view the derivation d on $\hat{\mathcal{A}}$ as an element of the Hochschild cohomology $HH^1(\hat{\mathcal{A}})$; two derivations on $\hat{\mathcal{A}}$ related by gauge change represent the same element of $HH^1(\hat{\mathcal{A}})$.

Finally, we define stabilization.

Definition 2.13. Let $(\hat{\mathcal{A}}, d)$ be an LSFT algebra generated by $q_1, \dots, q_n, p_1, \dots, p_n, t, t^{-1}$.

The degree- i (*algebraic*) *stabilization* of $(\hat{\mathcal{A}}, d)$ is the LSFT algebra $(S_i\hat{\mathcal{A}}, d)$ generated by $q_1, \dots, q_n, p_1, \dots, p_n, t, t^{-1}, q_a, q_b, p_a, p_b$, where q_a, q_b, p_a, p_b are four new generators with $|q_a| = |q_b| + 1 = -1 - |p_a| = -|p_b| = i$, and d is defined on $S_i\hat{\mathcal{A}}$ by extending the existing derivation by

$$d(q_a) = q_b, \quad d(q_b) = [F_d, q_a], \quad d(p_b) = p_a, \quad d(p_a) = [F_d, p_b].$$

If $(S_i\hat{\mathcal{A}}, d)$ is a stabilization of $(\hat{\mathcal{A}}, d)$, then we say that $(\hat{\mathcal{A}}, d)$ is a destabilization of $(S_i\hat{\mathcal{A}}, d)$.

On $\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}$, this definition reduces to Chekanov's notion of stabilization for DGAs. The following definition then generalizes Chekanov's stable tame isomorphism.

Definition 2.14. Two LSFT algebras are *equivalent* if they are related by some finite sequence of basis changes, gauge changes, stabilizations, and destabilizations.

Proposition 2.15. *An equivalence of LSFT algebras is a homotopy equivalence of filtered curved dg-algebras.*

Proof. Since basis changes and gauge changes are isomorphisms of the underlying algebra, it is easy to check that they are homotopy equivalences in the sense of Definition 2.4. It thus suffices to show that stabilization is a homotopy equivalence as well.

Let $(\hat{\mathcal{A}}, d)$ be an LSFT algebra with stabilization $(S_i\hat{\mathcal{A}}, d)$. Let $\iota : \hat{\mathcal{A}} \rightarrow S_i\hat{\mathcal{A}}$ and $\pi : S_i\hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ denote the usual inclusion and projection maps, where π projects away any word involving the four additional generators q_a, q_b, p_a, p_b . Then $(\iota, 0) : (\hat{\mathcal{A}}, d, F_d) \rightarrow (S_i\hat{\mathcal{A}}, d, F_d)$ and $(\pi, 0) : (S_i\hat{\mathcal{A}}, d, F_d) \rightarrow (\hat{\mathcal{A}}, d, F_d)$ are morphisms of filtered curved dg-algebras, and it is clear that $(\pi, 0) \circ (\iota, 0) = (\text{Id}, 0)$.

As for $(\iota, 0) \circ (\pi, 0)$, define a \mathbb{Z} -linear map $H : S_i\hat{\mathcal{A}} \rightarrow S_i\hat{\mathcal{A}}$ by its action on words w :

$$H(w) = \begin{cases} (-1)^{|w_1|} w_1 q_a w_2 & \text{if } w = w_1 q_b w_2 \text{ for words } w_1, w_2 \text{ with } w_1 \in \mathcal{A} \\ (-1)^{|w_1|} w_1 p_b w_2 & \text{if } w = w_1 p_a w_2 \text{ for words } w_1, w_2 \text{ with } w_1 \in \mathcal{A} \\ 0 & \text{if } w \in \mathcal{A} \text{ or } w = w_1 q_a w_2 \text{ or } w = w_1 p_b w_2 \text{ for} \\ & \text{words } w_1, w_2 \text{ with } w_1 \in \mathcal{A}. \end{cases}$$

The proof is complete once we check that H is a homotopy between the identity and $\iota \circ \pi$, a fact that we defer to the ensuing lemma. \square

Lemma 2.16. *On $S_i\hat{\mathcal{A}}$, we have $\text{Id}_{S_i\hat{\mathcal{A}}} - \iota \circ \pi = H \circ d + d \circ H$.*

Proof. It suffices to check

$$(2) \quad w - \iota \circ \pi(w) = (H \circ d)(w) + (d \circ H)(w)$$

for all words w in $S_i\hat{\mathcal{A}}$. If $w \in \mathcal{A}$, both sides of (2) are 0. Otherwise, the left hand side of (2) is w . If $w = w_1 q_a w_2$ for $w_1 \in \mathcal{A}$, then

$$(H \circ d)(w) + (d \circ H)(w) = (H \circ d)(w_1 q_a w_2) = w_1 q_b w_2 = w;$$

if $w = w_1 q_b w_2$ for $w_1 \in \mathcal{A}$, then

$$\begin{aligned} (H \circ d)(w) + (d \circ H)(w) &= H \left((dw_1) q_b w_2 + (-1)^{|w_1|} w_1 [F_d, q_a] w_2 \right. \\ &\quad \left. + (-1)^{|w_1|+i+1} w_1 q_b (dw_2) \right) + (-1)^{|w_1|} d(w_1 q_a w_2) \\ &= w. \end{aligned}$$

The cases $w = w_1 p_a w_2$ and $w = w_1 p_b w_2$ for $w_1 \in \mathcal{A}$ are similar. \square

2.3. Combinatorial description of the invariant. Let Λ be a Legendrian knot in \mathbb{R}^3 with the standard contact structure $\ker(dz - y dx)$, that is, a knot everywhere tangent to the contact structure. In this section, we associate an LSFT algebra to Λ . A generic knot Λ has finitely many Reeb chords R_1, \dots, R_n . To each Reeb chord R_j , we assign two indeterminates q_j, p_j . Let $\pi_{xy}(\Lambda)$ be the knot diagram given by projecting Λ to the xy plane;

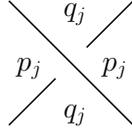


FIGURE 2.1. Labeling the quadrants at a crossing R_j in the xy projection $\pi_{xy}(\Lambda)$ by p_j and q_j .

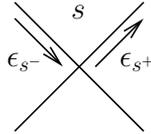


FIGURE 2.2. Signs at a quadrant s .

then the crossings of $\pi_{xy}(\Lambda)$ are the Reeb chords of Λ , and the four quadrants at each crossing can be labeled with a q or a p as shown in Figure 2.1. We also fix two points $*, \bullet$ on Λ , neither of which lies at an endpoint of a Reeb chord; the LSFT algebra will depend on the choices of $*, \bullet$, though the equivalence class of the LSFT algebra will not.

Recall that Λ has two classical invariants tb and r . The Thurston–Bennequin number $tb(\Lambda)$ is the writhe of the knot diagram $\pi_{xy}(\Lambda)$. The rotation number $r(\Lambda)$ is the Whitney index of $\pi_{xy}(\Lambda)$. More precisely, if $\gamma: [a, b] \rightarrow \mathbb{R}^2$ is any immersed path, then define $r(\gamma) \in \mathbb{R}$ to be the number of counterclockwise revolutions made by the unit tangent vector $\gamma'(t)/|\gamma'(t)|$ around S^1 as t goes from a to b ; $\pi_{xy}(\Lambda)$ is a closed immersed path and $r(\Lambda) = r(\pi_{xy}(\Lambda)) \in \mathbb{Z}$.

We now construct the LSFT algebra $\hat{\mathcal{A}}$ associated to $(\Lambda, *, \bullet)$. This is generated by $q_1, \dots, q_n, p_1, \dots, p_n, t, t^{-1}$, with grading as follows. For each $j = 1, \dots, n$, there is a unique path γ_j along $\pi_{xy}(\Lambda)$ beginning at the overcrossing of crossing R_j , ending at the undercrossing of R_j , and not passing through $*$. If we assume the crossings of $\pi_{xy}(\Lambda)$ are transverse, then $r(\gamma_j)$ is neither an integer nor a half-integer. Define

$$\begin{aligned} |q_j| &= \lfloor 2r(\gamma_j) \rfloor \\ |p_j| &= \lfloor -2r(\gamma_j) \rfloor = -1 - |q_j| \\ |t| &= -2r(\Lambda) \\ |t^{-1}| &= 2r(\Lambda). \end{aligned}$$

We note that the gradings for the q 's and t are the same as in Legendrian contact homology.

When considering signs in the theory, we will often draw an arrow alongside a section of Λ ; such an arrow is understood to correspond to a sign ± 1 , namely $+1$ if the arrow agrees with the given orientation of Λ , and -1 if it disagrees. In this vein, we have the following easy result.

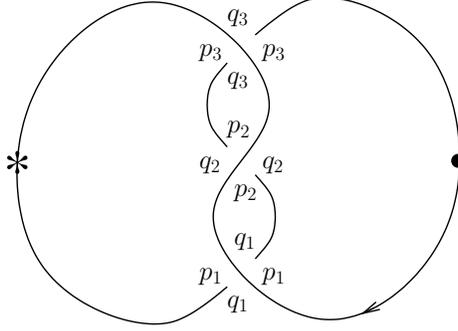


FIGURE 2.3. The xy projection of a Legendrian unknot Λ_0 . Note that this is the usual Legendrian unknot after a Reidemeister II move.

Lemma 2.17. *Let s be a p or q , corresponding to a corner at a crossing of $\pi_{xy}(\Lambda)$. Define the signs $\epsilon_{s-}, \epsilon_{s+} \in \{\pm 1\}$ to be the orientations along the sides of s relative to the orientation of Λ , as shown in Figure 2.2. Then $(-1)^{|s|} = \epsilon_{s-} \epsilon_{s+}$.*

In the language of [ENS02], s is “coherent” ($\epsilon_{s-} \epsilon_{s+} = 1$) if and only if $|s|$ is even.

Example. Through this section and Section 3, we will use the Legendrian knot Λ_0 depicted in Figure 2.3 as a running example. Here the gradings are given by

$$|q_2| = 2, \quad |q_1| = |q_3| = 1, \quad |t| = 0, \quad |p_1| = |p_3| = -2, \quad |p_2| = -3.$$

This agrees with the fact that p_1, q_2, p_3 are coherent while q_1, p_2, q_3 are not.

We define the derivation d on $\hat{\mathcal{A}}$ as the sum of two derivations $d_{\text{SFT}} + d_{\text{str}}$, where d_{SFT} is the “SFT differential” and d_{str} is the “string differential”.

For any $k \geq 1$, let D_k^2 denote the unit disk $\{|z| \leq 1\} \subset \mathbb{C}$ minus k fixed points $*_1, \dots, *_k$ on the boundary, ordered sequentially in counterclockwise order. The punctures divide the boundary ∂D_k^2 into k arcs denoted by $(\partial D_k^2)_1, \dots, (\partial D_k^2)_k$, where $(\partial D_k^2)_i$ is the portion of ∂D_k^2 between $*_i$ and $*_{i+1}$ (or between $*_k$ and $*_1$ if $i = k$).

Definition 2.18. For any s_1, \dots, s_k where $k \geq 1$ and each s is a q or p , let $\Delta(s_1, \dots, s_k)$ denote the set of all orientation-preserving immersions $f: (D_k^2, \partial D_k^2) \rightarrow (\mathbb{R}^2, \Lambda)$, up to domain reparametrization, such that $f(\partial D_k^2) \subset \Lambda$ and f sends neighborhoods of the boundary punctures $*_1, \dots, *_k$ to quadrants labeled s_1, \dots, s_k in succession.

We will call the quadrants described in Definition 2.18, labeled by s_1, \dots, s_k , the *corners* of f . Note that $\Delta(s_1, \dots, s_k)$ is unchanged by cyclic permutation of the s 's. We also have the following “index formula”.

Lemma 2.19. *Suppose that $f \in \Delta(s_1, \dots, s_k)$, and let α be the number of times $f|_{(\partial D_k^2)}$ passes through $*$, counted according to the orientation of Λ . Then*

$$|s_1| + \dots + |s_k| - 2\alpha r(\Lambda) = -2.$$

Proof. For each s_j , define γ_{s_j} to be the path in $\pi_{xy}(\Lambda)$ given by γ_{k_j} if $s_j = q_{k_j}$ and $-\gamma_{k_j}$ (i.e., γ_{k_j} with the opposite orientation) if $s_j = p_{k_j}$. Also define $\gamma_{f,j}$ to be the image in $\pi_{xy}(\Lambda)$ of $f|_{(\partial D_k^2)_j}$. Then

$$\gamma_{s_1} \cup \gamma_{f,1} \cup \gamma_{s_2} \cup \gamma_{f,2} \cup \dots \cup \gamma_{s_k} \cup \gamma_{f,k}$$

represents a closed loop in $\pi_{xy}(\Lambda)$ (more precisely, the projection of a closed loop in Λ) wrapping around $\pi_{xy}(\Lambda)$ α times. It follows that

$$\sum_{j=1}^k r(\gamma_{s_j}) + \sum_{j=1}^k r(\gamma_{f,j}) = \alpha r(\Lambda).$$

Now if θ_j is the angle (between 0 and π) determined by the image of f at $*_j$, then $r(\gamma_{s_j}) = n/2 - \theta_j/(2\pi)$ for some integer n , and thus $|s_j| = 2r(\gamma_{s_j}) - \theta_j/\pi$. On the other hand, since f is an immersed disk, $\sum_{j=1}^k (r(\gamma_{f,j}) + \theta_j/(2\pi)) = 1$. It follows that

$$\sum |s_j| = \sum (2r(\gamma_{s_j}) - \theta_j/\pi) = 2\alpha r(\Lambda) - \sum (2r(\gamma_{f,j}) - \theta_j/\pi) = 2\alpha r(\Lambda) - 2,$$

as desired. \square

For each map $f \in \Delta(s_1, \dots, s_k)$, we can define a word $w(f; s_1) \in \mathcal{A}$ by

$$w(f; s_1) = t^{\alpha_1} s_2 t^{\alpha_2} s_3 \dots t^{\alpha_{k-1}} s_k t^{\alpha_k},$$

where α_i is the number of times $f|_{(\partial D_k^2)_i}$ passes through $*$, counted according to the orientation of Λ . We also associate a sign $\epsilon(f; s_1) \in \{\pm 1\}$ to f as follows. Each quadrant of a crossing of $\pi_{xy}(\Lambda)$ can be given an orientation sign according to Figure 2.4. For each of the k corners of f , we thus obtain a sign $\epsilon_i(f)$. Further define a sign $\epsilon'(f; s_1)$ to be $+1$ if the image of $f|_{[*_1, *_2]} \subset \pi_{xy}(\Lambda)$, oriented from $f(*_1)$ to $f(*_2)$, has the same orientation as $\pi_{xy}(\Lambda)$, and -1 if it has the opposite orientation. Finally, define

$$\epsilon(f; s_1) = \epsilon'(f; s_1) \epsilon_1(f) \dots \epsilon_k(f).$$

See Figure 2.5.

Example. Consider the bigon f in Figure 2.3 with corners at p_2 and q_3 , which can be considered as an element of $\Delta(p_2, q_3)$ and of $\Delta(q_3, p_2)$. The orientation signs of both corners are -1 . If we consider $f \in \Delta(p_2, q_3)$, then $\epsilon(f; p_2) = \epsilon'(f; p_2) = 1$; if we consider $f \in \Delta(q_3, p_2)$, then $\epsilon(f; q_3) = \epsilon'(f; q_3) = -1$.

The following observation will be useful in Section 3.3.

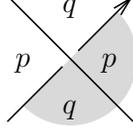


FIGURE 2.4. Orientation signs for corners. The two unshaded corners are given the sign $+1$, the shaded corners -1 . The arrow indicates the orientation of the knot and ensures that each crossing can be uniquely viewed as this local picture.

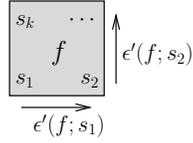


FIGURE 2.5. The immersed disk $f \in \Delta(s_1, \dots, s_k)$ contributes the term $\epsilon s_2 \dots s_k$ to $d_{\text{SFT}}(s_1^*)$, where ϵ is the product of: $\epsilon'(f; s_1)$ shown here; the orientation signs for the k corners s_1, \dots, s_k ; and $\{s_1, s_1^*\}$. The $\epsilon'(f; s_2)$ sign will be used in the proof of Lemma 3.10.

Lemma 2.20. *Any two diagonally-opposite corners at a crossing have opposite orientation signs. Also, if s, s^* denote consecutive corners at a crossing, and s lies counterclockwise from s^* , then the product of the orientation signs of s and s^* is $\{s, s^*\}$ (recall that this is $+1$ if s is a p , -1 if s is a q).*

Definition 2.21. Define the SFT differential on $\hat{\mathcal{A}}$ by

$$\begin{aligned} d_{\text{SFT}}(q_i) &= \sum_{f \in \Delta(p_i)} \epsilon(f; p_i) w(f) \\ d_{\text{SFT}}(p_i) &= - \sum_{f \in \Delta(q_i)} \epsilon(f; q_i) w(f) \\ d_{\text{SFT}}(t) &= d_{\text{SFT}}(t^{-1}) = 0, \end{aligned}$$

where $\Delta(s) = \coprod \Delta(s, s_2, \dots, s_k)$ is the set of all immersed disks with a corner at s (i.e., over all possible k and s_2, \dots, s_k). An immersed disk with multiple corners at s contributes multiple times to $d_{\text{SFT}}(s)$. Extend d_{SFT} to all of $\hat{\mathcal{A}}$ as a derivation.

It is possible that $d_{\text{SFT}}(q_i)$ or $d_{\text{SFT}}(p_i)$ may be an infinite sum, but it will always be a sum in the p -adic completion $\hat{\mathcal{A}}$; see the discussion of h in Section 3.

We note that d_{SFT} preserves the p filtration on $\hat{\mathcal{A}}$. This is a consequence of a basic area estimate originally due to Chekanov. Define a height function on

the p 's and q 's as follows: let $h(p_j)$ be the length of the Reeb chord R_j (i.e., the difference in the z coordinates of its endpoints), and let $h(q_j) = -h(p_j)$.

Lemma 2.22. *If $\Delta(s_1, \dots, s_k)$ is nonempty, then $\sum_{j=1}^k h(s_j) > 0$.*

Proof. Since $dz = y dx$ along Λ , it is easy to show from Stokes' Theorem that $\sum h(s_j)$ is the area of an immersed-disk element of $\Delta(s_1, \dots, s_k)$. See [Che02, ENS02]. \square

Lemma 2.23. *d_{SFT} has degree -1 and preserves the p filtration on $\hat{\mathcal{A}}$.*

Proof. The fact that d_{SFT} has degree -1 follows from Lemma 2.19. Since $h(q_j) < 0$ and $h(p_j) > 0$ for all j , Lemma 2.22 implies that any term in $d_{\text{SFT}}(p_j)$ must contain a p , and hence that d_{SFT} preserves the p filtration. \square

Example. For Λ_0 , we have

$$\begin{aligned} d_{\text{SFT}}(p_1) &= -p_2 \\ d_{\text{SFT}}(q_1) &= 1 + t + p_3 q_2 + q_2 p_3 t \\ d_{\text{SFT}}(p_2) &= -p_1 p_3 - p_3 t p_1 \\ d_{\text{SFT}}(q_2) &= -q_1 + q_3 \\ d_{\text{SFT}}(p_3) &= p_2 \\ d_{\text{SFT}}(q_3) &= 1 + t + q_2 p_1 + t p_1 q_2 \\ d_{\text{SFT}}(t) &= d_{\text{SFT}}(t^{-1}) = 0. \end{aligned}$$

Note that $d_{\text{SFT}}^2 \neq 0$, a fact that remains true even if we quotient by cyclic permutations of words. This is an example of the bubbling problem mentioned in the Introduction.

We next define the string differential d_{str} . For each Reeb chord R_j of Λ , write R_j^+, R_j^- for the endpoints of R_j , with the Reeb vector field flowing from R_j^- to R_j^+ (i.e., R_j^+ has the greater z coordinate). View q_j and p_j as the line segment R_j , oriented from R_j^- to R_j^+ for q_j and from R_j^+ to R_j^- for p_j . Let \mathcal{R} denote the set of Reeb-chord endpoints $\{R_1^+, R_1^-, \dots, R_n^+, R_n^-\}$.

Let Γ be the set of embedded paths $\gamma: [0, 1] \rightarrow \Lambda$ such that $\gamma^{-1}(\mathcal{R})$ is finite and $\gamma'(\tau) \neq 0$ whenever $\gamma(\tau) \in \mathcal{R}$. If $\gamma \in \Gamma$ and $\gamma(\tau) \in \mathcal{R}$, then we can define signs $\epsilon_1(\gamma; \tau), \epsilon_2(\gamma; \tau), \epsilon(\gamma; \tau)$ as follows: $\epsilon_1(\gamma; \tau)$ is $+1$ if $\gamma(\tau) = R_i^-$ and -1 if $\gamma(\tau) = R_i^+$; $\epsilon_2(\gamma; \tau)$ is the sign of the orientation of γ near τ , relative to the orientation of Λ there; and $\epsilon(\gamma; \tau) = \epsilon_1(\gamma; \tau)\epsilon_2(\gamma; \tau)$. Define a map $\delta: \Gamma \rightarrow \mathcal{A}$ as follows: for each τ such that $\gamma(\tau) \in \mathcal{R}$, define

$$\tilde{\delta}(\gamma; \tau) = \begin{cases} q_i p_i & \text{if } \gamma(\tau) = R_i^+ \\ p_i q_i & \text{if } \gamma(\tau) = R_i^-, \end{cases}$$

and then set

$$\delta(\gamma) = \sum_{\tau \in \gamma^{-1}(\mathcal{R}), \tau \neq 0, 1} \epsilon(\gamma; \tau) \tilde{\delta}(\gamma; \tau).$$

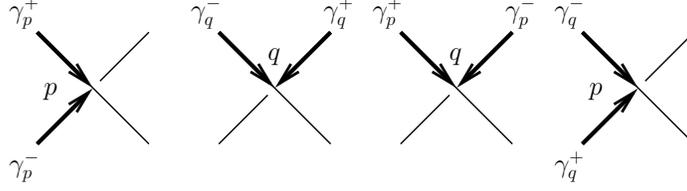


FIGURE 2.6. Four cases, left to right: p has holomorphic capping paths; q holomorphic; p antiholomorphic; q antiholomorphic. Note that the middle two diagrams are identical, as are the outer two.

We can now define the string portion d_{str} of the differential. Let s denote one of the q_i or p_i . Define s^+, s^- as follows: if $s = p_i$, then $s^\pm = R_i^\pm$; if $s = q_i$, then $s^\pm = R_i^\mp$. Recall that we are given two distinct points $*, \bullet \in \Lambda$. There are uniquely defined (up to reparametrization) injective paths γ_s^\pm in Λ that begin at \bullet , end at s^\pm , and do not pass through $*$. Note that $\gamma_s^+ = \gamma_{s^*}^-$ and $\gamma_s^- = \gamma_{s^*}^+$.

We distinguish two cases: if the quadrant at R_i in $\pi_{xy}(\Lambda)$ determined by the ends of γ_s^\pm is labeled by s , we say s has *holomorphic capping paths*; if it is labeled by s^* , we say s has *antiholomorphic capping paths*. Equivalently, s has holomorphic capping paths if and only if γ_s^- approaches the crossing in $\pi_{xy}(\Lambda)$ to the right of γ_s^+ . See Figure 2.6. Note that s has holomorphic capping paths if and only if s^* has antiholomorphic capping paths.

Definition 2.24. Define the string differential on $\hat{\mathcal{A}}$ as follows. If s is a p or q with holomorphic capping paths,

$$d_{\text{str}}(s) = \delta(\gamma_s^-) \cdot s + (-1)^{|s|} s \cdot \delta(\gamma_s^+);$$

if s is a p or q with antiholomorphic capping paths,

$$d_{\text{str}}(s) = \delta(\gamma_s^-) \cdot s + (-1)^{|s|} s \cdot \delta(\gamma_s^+) + \eta(s)s \cdot s^* \cdot s$$

where

$$\eta(s) = \begin{cases} +1 & \text{if } s = p \text{ and } \gamma_p^- \text{ is oriented like } \Lambda \\ -1 & \text{if } s = p \text{ and } \gamma_p^- \text{ is oriented unlike } \Lambda \\ -1 & \text{if } s = q \text{ and } \gamma_q^- \text{ is oriented like } \Lambda \\ +1 & \text{if } s = q \text{ and } \gamma_q^- \text{ is oriented unlike } \Lambda. \end{cases}$$

Furthermore, Λ itself can be viewed as a union of two injective paths $\gamma_{\Lambda,1}, \gamma_{\Lambda,2}$ where $\gamma_{\Lambda,1}$ begins at \bullet and ends at $*$, $\gamma_{\Lambda,2}$ begins at $*$ and ends at \bullet , and each path follows the orientation of Λ ; then set

$$d_{\text{str}}(t) = \delta(\gamma_{\Lambda,1}) \cdot t + t \cdot \delta(\gamma_{\Lambda,2})$$

and $d_{\text{str}}(t^{-1}) = -t^{-1} \cdot d_{\text{str}}(t) \cdot t^{-1}$. Extend d_{str} to all of $\hat{\mathcal{A}}$ as a derivation.

Note that d_{str} is well defined since $d_{\text{str}}(t \cdot t^{-1}) = d_{\text{str}}(t^{-1} \cdot t) = 0$.

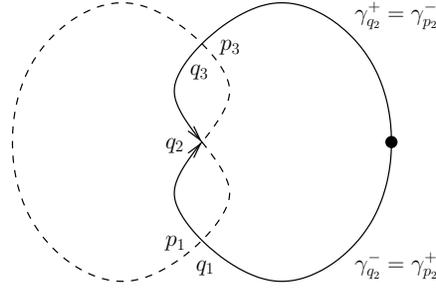


FIGURE 2.7. Capping paths $\gamma_{q_2}^\pm, \gamma_{p_2}^\pm$ of q_2, p_2 for the Legendrian knot from Figure 2.3.

Example. For Λ_0 , the capping paths for p_2, q_2 are depicted in Figure 2.7, leading to $d_{\text{str}}(q_2) = -q_1 p_1 q_2 + q_2 p_3 q_3$ and $d_{\text{str}}(p_2) = -p_3 q_3 p_2 - p_2 q_2 p_2 - p_2 q_1 p_1$. The full string differential is given below.

Theorem 2.25. $(\hat{\mathcal{A}}, d = d_{\text{SFT}} + d_{\text{str}})$ is an LSFT algebra.

The fact that d preserves the filtration on $\hat{\mathcal{A}}$ follows from the facts that d_{SFT} and d_{str} also preserve the filtration; this property for d_{SFT} has already been established, while for d_{str} this is clear by construction.

Theorem 2.25 is the LSFT analogue of the $d^2 = 0$ result in Legendrian contact homology, and indeed implies it. It will be proven in Section 3; see Proposition 3.15.

Example. For Λ_0 , the full derivation d is given by

$$\begin{aligned}
 d(p_1) &= -p_2 + (-p_3 q_3 p_1 - p_2 q_2 p_1) \\
 d(q_1) &= 1 + t + p_3 q_2 + q_2 p_3 t + (-q_1 p_1 q_1 - q_1 p_2 q_2 - q_1 p_3 q_3) \\
 d(p_2) &= -p_1 p_3 - p_3 t p_1 + (-p_3 q_3 p_2 - p_2 q_2 p_2 - p_2 q_1 p_1) \\
 d(q_2) &= -q_1 + q_3 + (-q_1 p_1 q_2 + q_2 p_3 q_3) \\
 d(p_3) &= p_2 + (p_3 q_2 p_2 + p_3 q_1 p_1) \\
 d(q_3) &= 1 + t + q_2 p_1 + t p_1 q_2 + (-q_1 p_1 q_3 - q_2 p_2 q_3 - q_3 p_3 q_3) \\
 d(t) &= (-q_1 p_1 t - q_2 p_2 t - q_3 p_3 t + t p_1 q_1 + t p_2 q_2 + t p_3 q_3),
 \end{aligned}$$

where the d_{str} contributions are enclosed in parentheses. The curvature for this differential is $F_d = -p_1 - p_3 - p_3 q_2 p_1$, and indeed it is straightforward to check that $d^2 s = [-p_1 - p_3 - p_3 q_2 p_1, s]$ for all generators s of the algebra, whence $d^2 x = [-p_1 - p_3 - p_3 q_2 p_1, x]$ for all $x \in \hat{\mathcal{A}}$.

Example. For reference and comparison, we give here the derivations for the standard Legendrian unknot and a once-stabilized Legendrian unknot with

$r = 1$, as shown in Figure 2.8. The former has

$$\begin{aligned} d(p) &= 0 \\ d(q) &= 1 + t - qpq \\ d(t) &= -qpt + tpq \end{aligned}$$

and $|p| = -2$, $|q| = 1$, $|t| = 0$, $F_d = -p$; the latter has

$$\begin{aligned} d(p_1) &= q_2 p_2 p_1 - p_1 p_2 q_2 \\ d(q_1) &= t - p_2 + p_2 q_2 q_1 - q_1 p_1 q_1 + q_1 q_2 p_2 \\ d(p_2) &= 0 \\ d(q_2) &= 1 - p_1 + q_2 p_2 q_2 \\ d(t) &= p_2 q_2 t - q_1 p_1 t + t p_1 q_1 - t q_2 p_2 \end{aligned}$$

and $|p_1| = 0$, $|q_1| = -1$, $|p_2| = -2$, $|q_2| = 1$, $|t| = -2$, $F_d = p_2$.

We next state the invariance result for LSFT algebras. Our invariance proof requires us to restrict to a special class of Legendrian isotopies, though we will see that this restriction covers all Legendrian isotopies if we instead restrict to particular types of xy projections.

Definition 2.26. Two xy projections Λ_1, Λ_2 of Legendrian knots are related by a *restricted Reidemeister II move* if there is an embedded disk $D \subset \mathbb{R}^3$ such that Λ_1, Λ_2 are identical outside D , each with exactly one crossing outside D , $\Lambda_1 \cap \partial D = \Lambda_2 \cap \partial D$ consists of two points, and $\Lambda_1 \cap D, \Lambda_2 \cap D$ are related by a Reidemeister II move inside D . See Figure 2.9.

Two xy projections are related by *restricted Reidemeister moves* if they are related by a sequence of Reidemeister III moves and restricted Reidemeister II moves; a *restricted Legendrian isotopy* is a Legendrian isotopy given in the xy projection by restricted Reidemeister moves.

Note that the Legendrian knot in Figure 2.3 is related to the standard one-crossing unknot by a Reidemeister II move eliminating crossings 2 and 3, but this move is not a restricted Reidemeister II move. It is clear that the knot from Figure 2.3 is not related to the standard unknot by restricted Reidemeister moves, though one can show that its LSFT algebra is equivalent to that of the standard unknot.

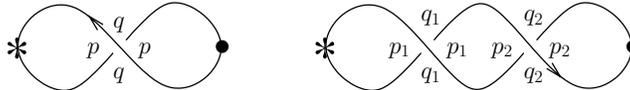


FIGURE 2.8. Two Legendrian unknots, the standard one with $tb = -1$ and $r = 0$ (left) and a once-stabilized one with $tb = -2$ and $r = 1$ (right).

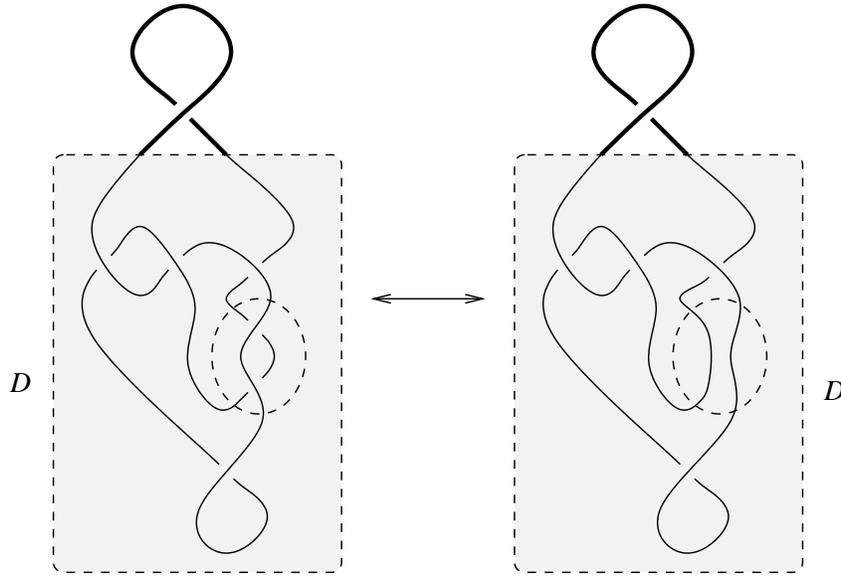


FIGURE 2.9. Restricted Reidemeister II move.

Recall that there is a standard procedure, called “morsification” [Fer02] or “resolution” [Ng03], to obtain an xy projection from a front (xz) projection of a Legendrian knot, by smoothing out left cusps and replacing right cusps by loops.

Proposition 2.27. *The resolutions of the fronts of two Legendrian isotopic knots can be related by restricted Reidemeister moves.*

Proof. Examine the resolutions of Legendrian Reidemeister moves for fronts: front Reidemeister III resolves to a usual Reidemeister III move; front Reidemeister I and II both resolve to Reidemeister II moves that are restricted since they do not involve the rightmost cusp of the front. \square

The next result is the LSFT version of invariance, and again implies the analogous result in contact homology.

Theorem 2.28. *If Λ and Λ' are related by restricted Legendrian isotopy, then the LSFT algebras for Λ and Λ' are equivalent.*

Theorem 2.28 will be proved in Section 4.

Corollary 2.29. *The LSFT algebra associated to the resolution of a Legendrian front is an invariant of the corresponding Legendrian knot.*

As mentioned in the Introduction, it is not unreasonable to guess that one can extend Theorem 2.28 to cover all Legendrian isotopies and not just restricted ones, but one might need to broaden the notion of equivalence to allow arbitrary curved A_∞ morphisms.

Corollary 2.30. *The stable tame isomorphism type of the contact homology DGA $(\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}, d)$ is invariant under restricted Legendrian isotopy.*

In fact, an examination of the proof of Theorem 2.28 shows that the contact homology DGA is invariant under all Legendrian isotopies, not just restricted ones; this recovers the original result of [Che02].

Remark 2.31 (The LSFT algebra and Legendrian contact homology). In Remark 2.7, we identified $\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}$ with the Chekanov–Eliashberg differential graded algebra [Che02, Eli98] calculating Legendrian contact homology. This holds not only in Chekanov’s original formulation over $\mathbb{Z}/2$, but also over \mathbb{Z} in the formulation of [EES05a, EES07, ENS02]. There are, however, two caveats to this identification. First, the signs used here do not coincide precisely with the signs from [ENS02], though they do agree with another sign assignment for Legendrian contact homology given in [EES05b]. However, up to a basis change, all possible sign assignments are equivalent. The precise statement is given and proven in Appendix A.

Second, there is a base ring issue. In the standard formulation of Legendrian contact homology, the differential graded algebra is generated by Reeb chords (the q_j ’s) over the group ring $\mathbb{Z}[H_1(\Lambda)]$, which for knots is $\mathbb{Z}[t, t^{-1}]$. In particular, $t^{\pm 1}$ commutes with all of the q_j ’s. By contrast, $\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}$ is generated by the q_j ’s and also $t^{\pm 1}$, with $d(t^{\pm 1}) = 0$, and $t^{\pm 1}$ does not commute with the q_j ’s. We can think of the contact homology differential graded algebra as a quotient of $\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}$ by commutators involving $t^{\pm 1}$.

On the other hand, there is no obvious reason why, in formulating Legendrian contact homology, we should impose the relation that t commutes with the q_j ’s. One could reasonably define Legendrian contact homology (even in situations more general than knots in \mathbb{R}^3) without this relation. In our case, we would precisely recover $\mathcal{F}^0\hat{\mathcal{A}}/\mathcal{F}^1\hat{\mathcal{A}}$.

2.4. The cyclic and commutative complexes. We now discuss two quotient complexes that can be derived from the LSFT algebra or any curved dg-algebra. The cyclic complex has close relations to string topology and the geometric motivation for the LSFT algebra; see Section 3.³ The commutative complex may be useful from a computational standpoint, especially since it has a particularly simple formulation in the case of the LSFT algebra, as we discuss at the end of this section.

Definition 2.32. Let $(\hat{\mathcal{A}}, d, F)$ be a curved dg-algebra.

³Cyclic constructions are common in Symplectic Field Theory and related topics. See for instance [BEE].

- (1) Let \mathcal{I} be the submodule of $\hat{\mathcal{A}}$ generated (over \mathbb{Z} , not over $\hat{\mathcal{A}}$) by commutators⁴

$$\{[x, y] \mid x, y \in \hat{\mathcal{A}} \text{ and at least one of } x, y \in \mathcal{F}^1\hat{\mathcal{A}}\}.$$

The *cyclic complex* associated to $(\hat{\mathcal{A}}, d, F)$ is $(\hat{\mathcal{A}}^{\text{cyc}} = \hat{\mathcal{A}}/\mathcal{I}, d)$, where d is the induced differential on $\hat{\mathcal{A}}^{\text{cyc}}$.

- (2) Let \mathcal{J} be the subalgebra of $\hat{\mathcal{A}}$ generated (over $\hat{\mathcal{A}}$) by commutators $[x, y]$ for all $x, y \in \hat{\mathcal{A}}$. The *commutative complex* associated to $(\hat{\mathcal{A}}, d, F)$ is $(\hat{\mathcal{A}}^{\text{comm}} = \hat{\mathcal{A}}/\mathcal{J}, d)$, where d is the induced differential on $\hat{\mathcal{A}}^{\text{comm}}$.

The key point here is that $d^2 = 0$ on $\hat{\mathcal{A}}^{\text{cyc}}$ and $\hat{\mathcal{A}}^{\text{comm}}$, by the definition of curved dg-algebra.

When $\hat{\mathcal{A}}$ is a tensor algebra (as for the LSFT algebra), $\hat{\mathcal{A}}^{\text{cyc}}$ is generated by “cyclic words”, or words modulo cyclic permutations of the letters (for words in $\mathcal{F}^1\hat{\mathcal{A}}$). Note that $\hat{\mathcal{A}}^{\text{cyc}}$ is a \mathbb{Z} -module and not an algebra; it however still inherits the grading and filtration from $\hat{\mathcal{A}}$. By contrast, $\hat{\mathcal{A}}^{\text{comm}}$ is a (sign-)commutative algebra over \mathbb{Z} , the polynomial algebra generated by the generators of $\hat{\mathcal{A}}$. There are obvious quotient maps

$$(\hat{\mathcal{A}}, d, F) \twoheadrightarrow (\hat{\mathcal{A}}^{\text{cyc}}, d) \twoheadrightarrow (\hat{\mathcal{A}}^{\text{comm}}, d),$$

and the latter induces a map on homology.

We next show that the cyclic and commutative complexes associated to the LSFT algebra of a Legendrian knot are invariant. This is a direct consequence of the following result.

Proposition 2.33. *If $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}', d')$ are equivalent LSFT algebras, then their rational cyclic and commutative quotient complexes are filtered chain homotopy equivalent. In particular, they are quasi-isomorphic:*

$$H_*(\hat{\mathcal{A}}^{\text{cyc}} \otimes \mathbb{Q}, d) \cong H_*((\hat{\mathcal{A}}')^{\text{cyc}} \otimes \mathbb{Q}, d')$$

as filtered graded \mathbb{Q} -modules, and

$$H_*(\hat{\mathcal{A}}^{\text{comm}} \otimes \mathbb{Q}, d) \cong H_*((\hat{\mathcal{A}}')^{\text{comm}} \otimes \mathbb{Q}, d')$$

as filtered graded \mathbb{Q} -algebras.

The substance of the proof of Proposition 2.33, which we give below, is invariance under stabilization. Recall from the proof of Proposition 2.15 the maps ι, π between an LSFT algebra $(\hat{\mathcal{A}}, d)$ and its stabilization $(S_i\hat{\mathcal{A}}, d)$. Though the homotopy operator H from that proof does not descend from $S_i\hat{\mathcal{A}}$ to $(S_i\hat{\mathcal{A}})^{\text{cyc}}$, we can define a slight variant that serves as the corresponding homotopy operator for $(S_i\hat{\mathcal{A}})^{\text{cyc}}$. Let $h: S_i\hat{\mathcal{A}} \rightarrow (S_i\hat{\mathcal{A}})^{\text{cyc}}$ be the derivation

⁴The restriction that one of $x, y \in \mathcal{F}^1\hat{\mathcal{A}}$ is unnecessary for most purposes, but is needed for the theory to include full Legendrian contact homology, rather than a cyclic version, as a quotient.

defined by $h(q_b) = q_a$, $h(p_a) = p_b$, $h(q_a) = h(p_b) = 0$, $h(q_j) = h(p_j) = h(t^{\pm 1}) = 0$ for all j (besides a, b). For a word $w \in \hat{\mathcal{A}}$, let $\sigma(w)$ be the total number of occurrences of q_a, q_b, p_a, p_b in x . Now define $H^{\text{cyc}}: (S_i \hat{\mathcal{A}}) \otimes \mathbb{Q} \rightarrow (S_i \hat{\mathcal{A}}) \otimes \mathbb{Q}$ by

$$H^{\text{cyc}}(w) = \begin{cases} 0 & \text{if } \sigma(w) = 0 \\ \frac{1}{\sigma(w)} h(w) & \text{if } \sigma(w) > 0. \end{cases}$$

Lemma 2.34. *On $(S_i \hat{\mathcal{A}}) \otimes \mathbb{Q}$, we have $\text{Id}_{S_i \hat{\mathcal{A}}} - \iota \circ \pi = H^{\text{cyc}} \circ d + d \circ H^{\text{cyc}}$.*

Proof. It suffices to show that $(h \circ d + d \circ h)(w) = \sigma(w)w$ for all words w , since d preserves the number of occurrences of q_a, q_b, p_a, p_b . But both $h \circ d + d \circ h$ and the map (generated on words by) $w \mapsto \sigma(w)w$ are derivations, and they agree on the generators $q_a, q_b, p_a, p_b, q_j, p_j, t$. \square

Proof of Proposition 2.33. Let $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}', d')$ be equivalent LSFT algebras. We show that $(\hat{\mathcal{A}}^{\text{cyc}} \otimes \mathbb{Q}, d)$ and $((\hat{\mathcal{A}}')^{\text{cyc}} \otimes \mathbb{Q}, d')$ are filtered chain homotopy equivalent; the proof of the corresponding result for the commutative complexes is nearly identical. The result clearly holds if $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}', d')$ are related by a basis change or gauge change. Thus we may assume that $(\hat{\mathcal{A}}', d') = (S_i \hat{\mathcal{A}}, d)$ is a stabilization of $(\hat{\mathcal{A}}, d)$. In this case, the inclusion and projection maps ι, π between $\hat{\mathcal{A}}$ and $S_i \hat{\mathcal{A}}$ induce chain maps $\iota: \hat{\mathcal{A}}^{\text{cyc}} \rightarrow (S_i \hat{\mathcal{A}})^{\text{cyc}}$ and $\pi: (S_i \hat{\mathcal{A}})^{\text{cyc}} \rightarrow \hat{\mathcal{A}}^{\text{cyc}}$. Furthermore, $\pi \circ \iota = \text{Id}_{\hat{\mathcal{A}}^{\text{cyc}}}$, while $\iota \circ \pi$ is chain homotopic over \mathbb{Q} to $\text{Id}_{(S_i \hat{\mathcal{A}})^{\text{cyc}}}$ by Lemma 2.34. \square

Corollary 2.35. *The cyclic and commutative complexes associated to the LSFT algebra of a Legendrian knot are invariant, up to filtered chain homotopy equivalence, under restricted Legendrian isotopy.*

We remark that it can be checked that the powers of F , $F^n \in \hat{\mathcal{A}}$ for $n \geq 0$, descend to invariant classes in $H(\hat{\mathcal{A}}^{\text{cyc}} \otimes \mathbb{Q}, d)$. See [Pos93] for a fuller discussion, where these invariant classes are called ‘‘Chern classes’’.

To end this section, we observe that the commutative complex for an LSFT algebra has a rather simpler formulation than the full LSFT algebra. More precisely, on $\hat{\mathcal{A}}^{\text{comm}}$ we can still define $d = d_{\text{SFT}} + d_{\text{str}}$, with d_{SFT} defined as for $\hat{\mathcal{A}}$, but now d_{str} can be given as follows.

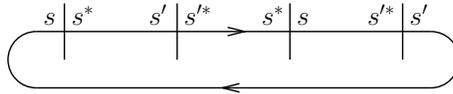


FIGURE 2.10. A schematic diagram traversing $\pi_{xy}(\Lambda)$ along its orientation. In this picture, the crossings corresponding to s and s' are interlaced, and we have $s \dashrightarrow s' \dashrightarrow s^* \dashrightarrow s'^* \dashrightarrow s$.

Definition 2.36. Two crossings in $\pi_{xy}(\Lambda)$ are *interlaced* if, in traversing the knot one full time, we encounter one crossing, then the other, then the first again, then the second again. If crossings corresponding to s and s' are interlaced, we say that s is *interlaced toward* s' and write $s \dashrightarrow s'$ if, when we traverse Λ along its orientation starting from s^- , we encounter $(s')^-$ before $(s')^+$.

See Figure 2.10 for an illustration. If the crossings corresponding to s and s' are interlaced, then there are two possibilities, $s \dashrightarrow s' \dashrightarrow s^* \dashrightarrow s'^* \dashrightarrow s$ or $s \dashrightarrow s'^* \dashrightarrow s^* \dashrightarrow s' \dashrightarrow s$, depending on the orientation of Λ .

Proposition 2.37. *In $\hat{\mathcal{A}}^{\text{comm}}$, we have $d_{\text{str}}(t^{\pm 1}) = 0$, while if s is a p or q , then*

$$d_{\text{str}}(s) = \sum_{s \dashrightarrow s'} \{s'^*, s'\} s'^* s' s$$

where the sum is over all s' such that s is interlaced toward s' , and $\{s'^*, s'\}$ is 1 if s' is a q , -1 if s' is a p .

The proof of Proposition 2.37 is an exercise in chasing signs, and we leave it to the interested reader. Note that d_{str} does not depend on the choice of base points $\bullet, *$ on Λ . Thus the differential d on $\hat{\mathcal{A}}^{\text{comm}}$ is independent of \bullet , and only depends on $*$ insofar as $*$ keeps track of powers of t in the SFT differential, cf. group-ring coefficients in Legendrian contact homology [ENS02].

3. STRING INTERPRETATION OF THE LSFT ALGEBRA

It will be useful to have another description of the LSFT algebra, closer to the standard SFT formalism and string topology. This allows us to prove the “ $d^2 = 0$ result”, Theorem 2.25.

3.1. Broken closed strings and the SFT bracket. Generators of the LSFT algebra of a Legendrian knot Λ are more conveniently seen as strings on Λ . From the holomorphic perspective, these are the boundaries of holomorphic disks with boundary on $\Lambda \times \mathbb{R}$. It will be fruitful, however, to consider all possible strings, not just those that arise as the boundary of a disk.

Definition 3.1. Let $\Lambda \in \mathbb{R}^3$ be a Legendrian knot with Reeb chords R_1, \dots, R_n , and let $R_i^\pm \in \Lambda$ denote the endpoints of Reeb chord R_i . For fixed $k \geq 0$, choose k distinct points τ_1, \dots, τ_k on an oriented circle S^1 so that they appear sequentially in order; we refer to these points as *punctures* of S^1 , and the punctures divide S^1 into k intervals, which we denote $[\tau_1, \tau_2], [\tau_2, \tau_3], \dots, [\tau_k, \tau_1]$. A *broken closed string* of length k is a piecewise continuous map $\gamma: S^1 \rightarrow \Lambda$ such that:

- (1) $\gamma|_{[\tau_j, \tau_{j+1}]}$ is continuous for each j ;
- (2) for each $j = 1, \dots, k$, either $\lim_{\tau \rightarrow \tau_j^\pm} \gamma(\tau) = R_{i_j}^\pm$ or $\lim_{\tau \rightarrow \tau_j^\pm} \gamma(\tau) = R_{i_j}^\mp$ for some i_j .

We consider broken closed strings up to orientation-preserving reparametrization of the domain S^1 .

If we are given a point $\bullet \in \Lambda$ distinct from any of the R_i^\pm , then choose a point $\tau_0 \in (\tau_k, \tau_1) \subset S^1$; a *based broken closed string* of length k is a broken closed string γ of length k such that $\gamma(\tau_0) = \bullet$.

Given distinct points $*, \bullet \in \Lambda$, we obtain a map w between based broken closed strings and words in $\hat{\mathcal{A}}$. If γ is a based broken closed string of length k , then define the word associated to γ to be

$$w(\gamma) = t^{a_0} s_1 t^{a_1} s_2 \cdots s_k t^{a_k}$$

where t^{a_j} is the number of times $\gamma|_{[\tau_j, \tau_{j+1}]}$ passes through $*$, counted with sign according to the orientation of Λ , and

$$s_j = \begin{cases} p_{i_j} & \text{if } \lim_{\tau \rightarrow \tau_j^\pm} \gamma(\tau) = R_{i_j}^\pm \\ q_{i_j} & \text{if } \lim_{\tau \rightarrow \tau_j^\pm} \gamma(\tau) = R_{i_j}^\mp. \end{cases}$$

Note that the correspondence between based broken closed strings and words in $\hat{\mathcal{A}}$ is bijective if we mod out the strings by homotopy.

Similarly, we can define a map, which we also denote by w , between broken closed strings and words in $\hat{\mathcal{A}}^{\text{cvc}}$. Note that this map does not depend on the choice of \bullet , as changing \bullet corresponds to conjugation by some power of t .

If γ, γ' are broken closed strings of length k, k' respectively, and a puncture from each is mapped to (the endpoints of) the same Reeb chord but in opposite directions, then we can glue γ, γ' at this puncture to obtain another broken closed string of length $k + k' - 2$. More precisely, suppose that the S^1 domain of γ has punctures τ_1, \dots, τ_k , the S^1 domain of γ' has punctures $\tau'_1, \dots, \tau'_{k'}$, and there are j and j' such that

$$\lim_{\tau \rightarrow \tau_j^\pm} \gamma(\tau) = \lim_{\tau \rightarrow \tau'_{j'}^\mp} \gamma'(\tau);$$

then we can define a broken closed string $\gamma * \gamma'$ on S^1 with sequential punctures $\tau_{j+1}, \dots, \tau_k, \tau_1, \dots, \tau_{j-1}, \tau'_{j'+1}, \dots, \tau'_{k'}, \tau'_1, \dots, \tau'_{j'-1}$ by

$$(\gamma * \gamma')|_{[\tau_i, \tau_{i+1}]} = \gamma|_{[\tau_i, \tau_{i+1}]}$$

for all $i = 1, \dots, k$ with $i \neq j$,

$$(\gamma * \gamma')|_{[\tau'_i, \tau'_{i+1}]} = \gamma'|_{[\tau'_i, \tau'_{i+1}]}$$

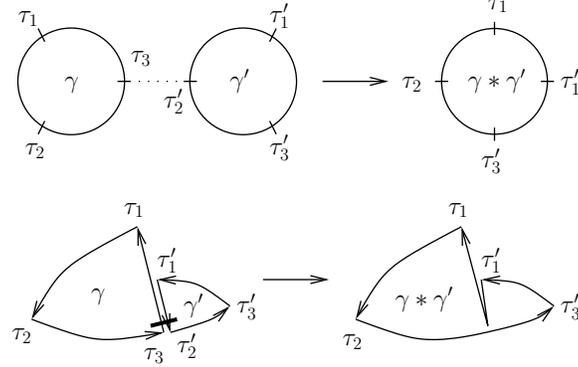


FIGURE 3.1. Gluing broken closed strings γ, γ' at punctures τ_3, τ'_2 to produce a new broken closed string $\gamma * \gamma'$. The top illustrates the picture in the domain; the bottom illustrates the image of $\gamma, \gamma', \gamma * \gamma'$ in the xy projection $\pi_{xy}(\Lambda)$. The heavy bar indicates the glued corners.

for all $i = 1, \dots, k'$ with $i \neq j'$,

$$\begin{aligned} (\gamma * \gamma')|_{[\tau_{j-1}, \tau'_{j'+1}]} &= \gamma|_{[\tau_{j-1}, \tau_j]} \cup \gamma'|_{[\tau'_{j'}, \tau'_{j'+1}]}, \\ (\gamma * \gamma')|_{[\tau'_{j'-1}, \tau_{j+1}]} &= \gamma'|_{[\tau'_{j'-1}, \tau'_{j'}]} \cup \gamma|_{[\tau_j, \tau_{j+1}]}. \end{aligned}$$

See Figure 3.1 for an illustration.

Using the gluing operation, we can define the SFT bracket of two broken closed strings to be the sum of all possible gluings of the broken closed strings. This gives an operation $\{\cdot, \cdot\}: \hat{\mathcal{A}}^{\text{cyc}} \otimes \hat{\mathcal{A}}^{\text{cyc}} \rightarrow \hat{\mathcal{A}}^{\text{cyc}}$ (at least mod 2). In the same way, we can define the SFT bracket of a broken closed string and a based broken closed string to be the sum of all based broken closed strings obtained by gluing; this gives a mod 2 map $\{\cdot, \cdot\}: \hat{\mathcal{A}}^{\text{cyc}} \otimes \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$. We refer to either operation as the *SFT bracket*. See Figure 3.2.

We can define the SFT bracket in a more precise algebraic manner, with the added benefit of lifting to \mathbb{Z} , as follows. First, given a word $w \in \mathcal{A}$ ending in a p or q , define a contraction map $\iota_w: \mathcal{A} \rightarrow \mathcal{A}$ as follows. Write $w = w'p_i$ or $w = w'q_i$ for some word w' , and set

$$\iota_{w'p_i}(s) = \begin{cases} \delta_{ij}w' & \text{if } x = q_j \\ 0 & \text{if } x = p_j \text{ or } x = t^{\pm 1} \end{cases}$$

and

$$\iota_{w'q_i}(x) = \begin{cases} -\delta_{ij}w' & \text{if } x = p_j \\ 0 & \text{if } x = q_j \text{ or } x = t^{\pm 1}, \end{cases}$$

where δ_{ij} is the Kronecker delta function; extend ι_w to a map $\mathcal{A} \rightarrow \mathcal{A}$ by linearity and the following modified Leibniz rule:

$$\iota_w(xy) = (\iota_w x)y + (-1)^{(|w|+1)|x|}x(\iota_w y).$$

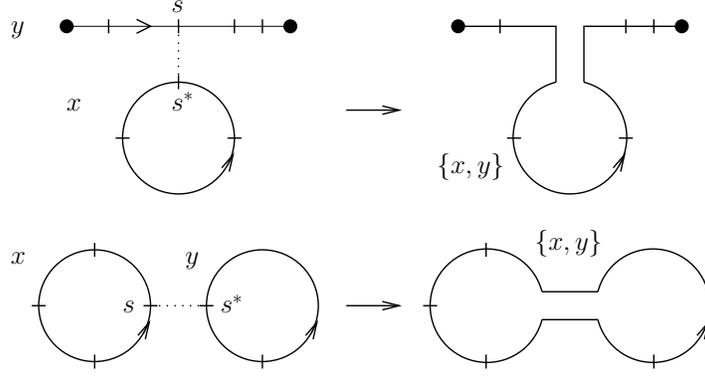


FIGURE 3.2. Gluing $x \in \hat{\mathcal{A}}^{\text{cyc}}$ to $y \in \hat{\mathcal{A}}$ to get $\{x, y\} \in \hat{\mathcal{A}}$ (top); gluing $x \in \hat{\mathcal{A}}^{\text{cyc}}$ to $y \in \hat{\mathcal{A}}^{\text{cyc}}$ to get $\{x, y\} \in \hat{\mathcal{A}}^{\text{cyc}}$ (bottom). The notches represent corners (punctures), and the dots represent the base point \bullet .

The unusual sign ensures that ι_w descends to a map $\mathcal{A}^{\text{cyc}} \rightarrow \mathcal{A}^{\text{cyc}}$.

Before extending contraction to cyclic words and defining the SFT bracket, we introduce notation for the set of all words that project to a particular cyclic word.

Definition 3.2. Let w be a word in \mathcal{A} . The *length* $l(w)$ of w is the number of q 's and p 's in w . The *cyclic word set* $\mathcal{C}(w)$ of w is the $l(w)$ -element multiset in $\hat{\mathcal{A}}$ of words equal to w in $\hat{\mathcal{A}}^{\text{cyc}}$. More precisely, if $w = t^{a_0} s_1 t^{a_1} s_2 \cdots s_{l(w)} t^{a_{l(w)}}$, where each s_i is a q or p , then

$$\begin{aligned} \mathcal{C}(w) = & \left\{ t^{a_0+a_{l(w)}} s_1 t^{a_1} s_2 \cdots s_{l(w)}, \right. \\ & (-1)^{|s_1|(|w|-|s_1|)} t^{a_1} s_2 \cdots s_{l(w)} t^{a_0+a_{l(w)}} s_1, \dots, \\ & \left. (-1)^{|s_{l(w)}|(|w|-|s_{l(w)})|} t^{a_{l(w)-1}} s_{l(w)} t^{a_0+a_{l(w)}} s_1 \cdots s_{l(w)-1} \right\}. \end{aligned}$$

Note that if w, w' represent the same element in $\hat{\mathcal{A}}^{\text{cyc}}$, then $\mathcal{C}(w) = \mathcal{C}(w')$.

Now if w is a word in \mathcal{A} and $[w]$ is the image of w in \mathcal{A}^{cyc} , then we define a contraction map $\iota_{[w]}: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ by

$$\iota_{[w]}(x) = \sum_{y \in \mathcal{C}(w)} \iota_y(x).$$

Here we use the convention that $\iota_{-w}(x) = -\iota_w(x)$ if w is a word in \mathcal{A} . The contraction map extends by linearity to a map $\iota(\cdot): \hat{\mathcal{A}}^{\text{cyc}} \otimes \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$.

Definition 3.3. The *SFT bracket* $\{\cdot, \cdot\}: \hat{\mathcal{A}}^{\text{cyc}} \otimes \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ is defined by $\{x, y\} = \iota_x(y)$. This descends to a map $\hat{\mathcal{A}}^{\text{cyc}} \otimes \hat{\mathcal{A}}^{\text{cyc}} \rightarrow \hat{\mathcal{A}}^{\text{cyc}}$, which we also denote by $\{\cdot, \cdot\}$.

For reference, the full sign rule, which can be deduced from the definition of ι , is as follows: the s and s^* entries in two words $w_1 s w_2$, $w_3 s^* w_4$ pair together to give

$$(3) \{w_1 s w_2, w_3 s^* w_4\} = (-1)^{|w_2||w_1 s| + (|w_1 s w_2| + 1)|w_3|} \{s, s^*\} w_3 w_2 w_1 w_4 + \cdots .$$

Proposition 3.4 (Properties of the SFT bracket). (1) Let $x, y \in \hat{\mathcal{A}}^{\text{cyc}}$.

Then

$$\{y, x\} = (-1)^{|x||y| + |x| + |y|} \{x, y\}.$$

(2) Let $x \in \hat{\mathcal{A}}^{\text{cyc}}$ and $y, z \in \hat{\mathcal{A}}$. Then

$$\{x, yz\} = \{x, y\}z + (-1)^{(|x| + 1)|y|} y\{x, z\}.$$

(3) Let $x, y \in \hat{\mathcal{A}}^{\text{cyc}}$ and $z \in \hat{\mathcal{A}}$. Then we have the following version of the Jacobi identity:

$$\{x, \{y, z\}\} + (-1)^{|x||y| + |x| + |y|} \{y, \{x, z\}\} = \{\{x, y\}, z\}.$$

Proof. We first establish the proposition mod 2. Note that (1) is clear, while (2) can be pictorially represented:

$$\{x, yz\} = \frac{yz}{\textcircled{x}} = \frac{y}{\textcircled{x}} \frac{z}{\textcircled{x}} + \frac{y}{\textcircled{x}} \frac{z}{\textcircled{x}} = \{x, y\}z + y\{x, z\}.$$

For (3), we have

$$\begin{aligned} \{x, \{y, z\}\} &= \frac{z}{\textcircled{x} \textcircled{y}} + \frac{z}{\textcircled{y} \textcircled{x}} + \frac{z}{\textcircled{x} \textcircled{y}} \\ \{y, \{x, z\}\} &= \frac{z}{\textcircled{x} \textcircled{y}} + \frac{z}{\textcircled{y} \textcircled{x}} + \frac{z}{\textcircled{x} \textcircled{y}} \\ \{\{x, y\}, z\} &= \frac{z}{\textcircled{x} \textcircled{y}} + \frac{z}{\textcircled{x} \textcircled{y}}. \end{aligned}$$

Checking the signs is now a routine exercise using equation (3). \square

3.2. The δ map. Having defined the SFT bracket, we now define another operation on strings, the δ map. This is essentially a string cobracket operation in the language of string topology. First we need to take a slightly closer look at broken closed strings.

Definition 3.5. A *generic broken closed string* is a broken closed string $\gamma: (S^1; \tau_1, \dots, \tau_k) \rightarrow \Lambda$ such that whenever $\gamma'(\tau) = 0$, $\gamma(\tau) \notin \mathcal{R}$, where we recall that \mathcal{R} is the set of Reeb chord endpoints; in particular, $\gamma'(\tau_i^\pm) \neq 0$, where $\gamma'(\tau_i^\pm) = \lim_{\tau \rightarrow \tau_i^\pm} \gamma'(\tau)$.

A generic broken closed string *has holomorphic corners* if for each i , $(\gamma'(\tau_i^-), \gamma(\tau_i^+) - \gamma(\tau_i^-), \gamma'(\tau_i^+))$ is a positively oriented frame in \mathbb{R}^3 . This condition is most easily interpreted in the xy projection: the image of γ in $\pi_{xy}(\Lambda)$ near each τ_i makes a left turn at the corner.

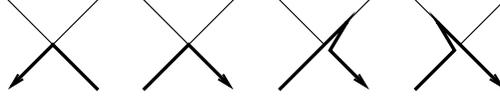


FIGURE 3.3. From left to right: a holomorphic corner of a broken closed string; a non-holomorphic corner; and two broken closed strings with holomorphic corners homotopic to the one with the non-holomorphic corner.

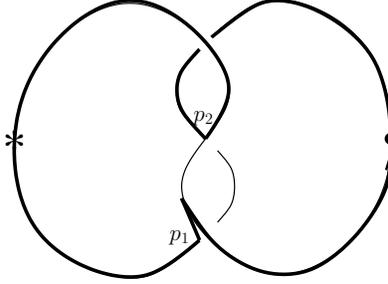


FIGURE 3.4. A generic broken closed string γ with holomorphic corners. The word associated to γ is p_2tp_1 .

See Figure 3.3 for an illustration of holomorphic corners and Figure 3.4 for an example of a generic broken closed string with holomorphic corners. It is easy to see that any broken closed string is homotopic to a generic broken closed string with holomorphic corners.

Now suppose that $\gamma: (S^1; \tau_1, \dots, \tau_k) \rightarrow \Lambda$ is a generic broken closed string of length k , and suppose $\tau \in (\tau_i, \tau_{i+1})$ satisfies $\gamma(\tau) = R_j^\pm$ for some j and some choice of \pm ; in this case, we say that τ is *interior Reeb* for γ . We can then define a broken closed string $\delta(\gamma; \tau)$ of length $k + 2$ to have sequential punctures $\tau_1, \dots, \tau_i, \tau^{(1)}, \tau^{(2)}, \tau_{i+1}, \dots, \tau_k$, and

$$\begin{aligned} \delta(\gamma; \tau)|_{[\tau_i, \tau^{(1)}]} &= \gamma|_{[\tau_i, t]} \\ \delta(\gamma; \tau)|_{[\tau^{(1)}, \tau^{(2)}]} &= \text{constant path at } R_j^\mp \\ \delta(\gamma; \tau)|_{[\tau^{(2)}, \tau_{i+1}]} &= \gamma|_{[\tau, \tau_{i+1}]} \\ \delta(\gamma; \tau)|_{[\tau_j, \tau_{j+1}]} &= \gamma|_{[\tau_j, \tau_{j+1}]} \quad \text{for } j \neq i. \end{aligned}$$

Note that $\delta(\gamma; \tau)$ is not generic, but it can be perturbed to become generic; furthermore, if we stipulate that the perturbed broken closed string has holomorphic corners on the domain interval (τ_i, τ_{i+1}) , then the perturbation is unique up to homotopy through generic broken closed strings with holomorphic corners.

If γ is a generic based broken closed curve and τ is interior Reeb for γ , we can similarly define a based broken closed curve $\delta(\gamma; \tau)$.

We can now define a map $\delta: \mathcal{A} \rightarrow \mathcal{A}$. In fact, δ is just a string reformulation of d_{str} from Section 2.3; see Proposition 3.9 below.

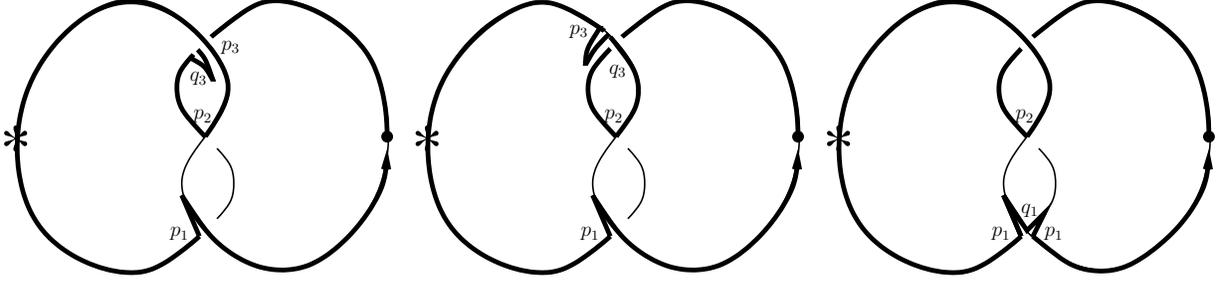


FIGURE 3.5. The three terms in $\delta(p_2tp_1)$ for the broken closed string $\gamma = \gamma(p_2tp_1)$ from Figure 3.4. Each new broken closed string has been perturbed to have holomorphic corners.

Definition 3.6. Let w be a word in \mathcal{A} ; we have $w = w(\gamma)$ for some generic based broken closed string with holomorphic corners. Define $\delta(w) \in \mathcal{A}$ by

$$\delta(w) = \sum_{\tau \text{ interior Reeb for } \gamma} \epsilon(\gamma; \tau) w(\delta(\gamma; \tau)),$$

with $\epsilon(\gamma; \tau)$ a sign to be defined in the next paragraph. Extend δ to \mathcal{A} by linearity.

We define $\epsilon(\gamma; \tau) \in \{\pm 1\}$ as follows. Suppose that $w(\gamma)$, with powers of t omitted, is of the form $s_1 \cdots s_k$, where s_j corresponds to $\gamma(\tau_j)$, and that $\tau_i < \tau < \tau_{i+1}$. Let $\epsilon_1 = \pm 1$ according to whether $\gamma(\tau) = R_j^{\pm 1}$, and let $\epsilon_2 = \pm 1$ according to whether the orientation of γ in a neighborhood of τ agrees or disagrees with the orientation of Λ . Finally, define

$$\epsilon(\gamma; \tau) = (-1)^{|s_1 \cdots s_i|} \epsilon_1 \epsilon_2.$$

See Figure 3.5 for an illustration of Definition 3.6. In short, if $w = s_1 \cdots s_k$, then $\delta(w)$ is a sum of terms of the form

$$\epsilon_2 (-1)^{|s_1 \cdots s_i|} s_1 \cdots s_i (pq) s_{i+1} \cdots s_k \quad \text{and} \quad -\epsilon_2 (-1)^{|s_1 \cdots s_i|} s_1 \cdots s_i (qp) s_{i+1} \cdots s_k,$$

where ϵ_2 measures the orientation of the broken closed string for w at the point where pq or qp is attached.

Example. Consider the broken closed string from Figure 3.4, corresponding to the word p_2tp_1 . From Figure 3.5, we see that $\delta(p_2tp_1)$ has three terms corresponding to $(p_3q_3)p_2tp_1$, $p_2(q_3p_3)tp_1$, and $p_2tp_1(q_1p_1)$. In fact, we have

$$\delta(p_2tp_1) = -p_3q_3p_2tp_1 + p_2q_3p_3tp_1 - p_2tp_1q_1p_1.$$

We verify the sign of the third term as an example. There are three -1 's contributing to this sign, one from $(-1)^{|p_2tp_1|}$, one from the fact that q_1 precedes p_1 in the parenthesis, and one from the fact that at the point where q_1p_1 is added, the broken closed string is oriented opposite to the orientation of the knot from Figure 2.3.

We remark that it can be readily checked that $\delta(p_2tp_1)$ agrees with $d_{\text{str}}(p_2tp_1)$, calculated from the values for $d_{\text{str}}(p_1)$, $d_{\text{str}}(p_2)$, $d_{\text{str}}(t)$ given in Section 2.3; see Proposition 3.9 below.

We now prove some fundamental algebraic properties of δ .

Definition 3.7. Define a \mathbb{Z} -linear map $\bullet: \hat{\mathcal{A}}^{\text{cyc}} \rightarrow \hat{\mathcal{A}}$ as follows. Any word w in \mathcal{A} corresponds to a based broken closed string $\gamma_w: (S^1, \tau_0, \tau_1, \dots, \tau_k) \rightarrow \Lambda$, and γ_w can be chosen so that whenever $\gamma_w(\tau) = \bullet$, $\gamma'_w(\tau) \neq 0$. Then γ_w passes through \bullet some number of times $n(\gamma_w) \in \mathbb{Z}$, counted with sign according to the orientation of the knot. (More precisely, since γ_w begins and ends at \bullet , one should “close up” γ_w and view it as a homotopy class of unbased broken closed strings when calculating $n(\gamma_w)$.) Now define $\bullet: \mathcal{A} \rightarrow \mathcal{A}$ by

$$\bullet(w) = (-1)^{|w|} n(\gamma_w) w$$

and $\bullet: \hat{\mathcal{A}}^{\text{cyc}} \rightarrow \hat{\mathcal{A}}$ by $\bullet([w]) = \sum_{w' \in \mathcal{C}(w)} \bullet(w')$.

Proposition 3.8 (Properties of δ). (1) δ gives a well-defined map from \mathcal{A} to \mathcal{A} and induces a well-defined map from \mathcal{A}^{cyc} to \mathcal{A}^{cyc} , as well as from $\hat{\mathcal{A}}$ to $\hat{\mathcal{A}}$ and from $\hat{\mathcal{A}}^{\text{cyc}}$ to $\hat{\mathcal{A}}^{\text{cyc}}$.

- (2) If $x, y \in \mathcal{A}$, then $\delta(xy) = (\delta x)y + (-1)^{|x|} x(\delta y)$.
- (3) If $x \in \mathcal{A}$, then $\delta^2(x) = 0$.
- (4) If $x, y \in \mathcal{A}^{\text{cyc}}$, then

$$\delta\{x, y\} = \{\delta x, y\} - (-1)^{|x|} \{x, \delta y\}.$$

- (5) If $x \in \mathcal{A}^{\text{cyc}}$ and $y \in \mathcal{A}$, then

$$\delta\{x, y\} = \{\delta x, y\} - (-1)^{|x|} \{x, \delta y\} + [\bullet(x), y].$$

Proof. For (1), note that any two generic based broken closed strings with holomorphic corners that represent the same word in \mathcal{A} can be related by a set of local moves, depicted in Figure 3.6. It is easy to check that δ is unchanged by each of these moves, and (1) follows. Items (2) and (3) are both clear mod 2, and are readily seen to hold over \mathbb{Z} by the definition of the signs $\epsilon(\gamma; \tau)$.

It remains to prove (4) and (5). Assume x is a cyclic word ($x \in \mathcal{A}^{\text{cyc}}$) and y is either a cyclic word or a word ($y \in \mathcal{A}^{\text{cyc}}$ or $y \in \mathcal{A}$), and define

$$f(x, y) = \delta\{x, y\} - \{\delta x, y\} + (-1)^{|x|} \{x, \delta y\}.$$

Most terms in $\delta\{x, y\}$ have an obvious corresponding term (with the same sign) in one of $\{\delta x, y\}$ or $(-1)^{|x|+1} \{x, \delta y\}$, and conversely. The exceptions are terms where the δ operation interacts with the bracket:

- every term in $\{x, y\}$ arises from gluing a corner s in x to s^* in y ; the resulting broken closed string has a segment that passes over the s, s^* crossing, where δ can be applied;

- every term in δx includes two consecutive corners not appearing in x , and either can be glued to y ;
- every term in δy includes two consecutive corners not appearing in y , and either can be glued to x .

These “exceptional terms” in $f(x, y)$ can be depicted as in Figure 3.7, where in each schematic picture x and y are oriented counterclockwise. Note that in each case, a quadrant of x or δx is glued to a quadrant of y or δy . Figure 3.7 only shows the terms where the x quadrant lies counterclockwise from the y quadrant. There is an analogous set of exceptional terms where the x quadrant lies clockwise from the y quadrant. Furthermore, there is a one-to-one correspondence between “counterclockwise” terms and “clockwise” terms; see Figure 3.8. In $\hat{\mathcal{A}}^{\text{cyc}}$, the terms under the one-to-one correspondence cancel pairwise in $f(x, y)$, and (4) follows (up to sign, which will be more carefully considered below).

For (5), the same cancellation holds, but we need to examine the position of the base point on the based broken closed string y . If x and y do not overlap (share a segment), then $\bullet(x) = f(x, y) = 0$ since none of the exceptional terms exist.

Otherwise, assume for simplicity that x and y share exactly one segment, and in particular x passes through the base point exactly once; the argument is similar in the general case. It follows that there are exactly two exceptional terms contributing to $f(x, y)$, and they are paired under the correspondence of Figure 3.8. (The configuration of x and y determines which particular pair from Figure 3.8 appears.)

First work mod 2. If the base point does not lie on the shared segment between x and y , then $\bullet(x) = f(x, y) = 0$ since the two exceptional terms cancel in \mathcal{A} . If the base point lies on the shared segment, then one exceptional term contributes $(\bullet(x))(y)$ and the other $(y)(\bullet(x))$, because the base point is positioned differently on the glued broken closed string depending

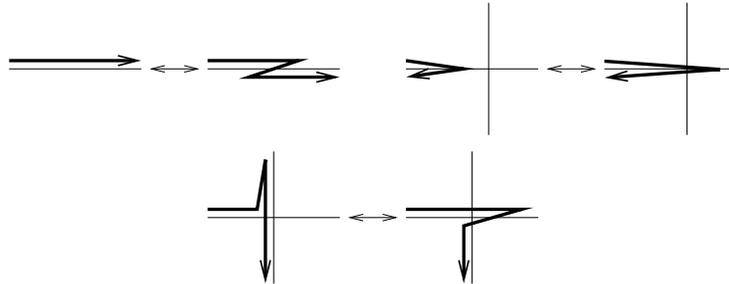


FIGURE 3.6. Local moves in the xy projection relating two broken closed strings that give the same word in \mathcal{A} . The broken closed strings are drawn thickly and the underlying xy projection of Λ is drawn thinly.

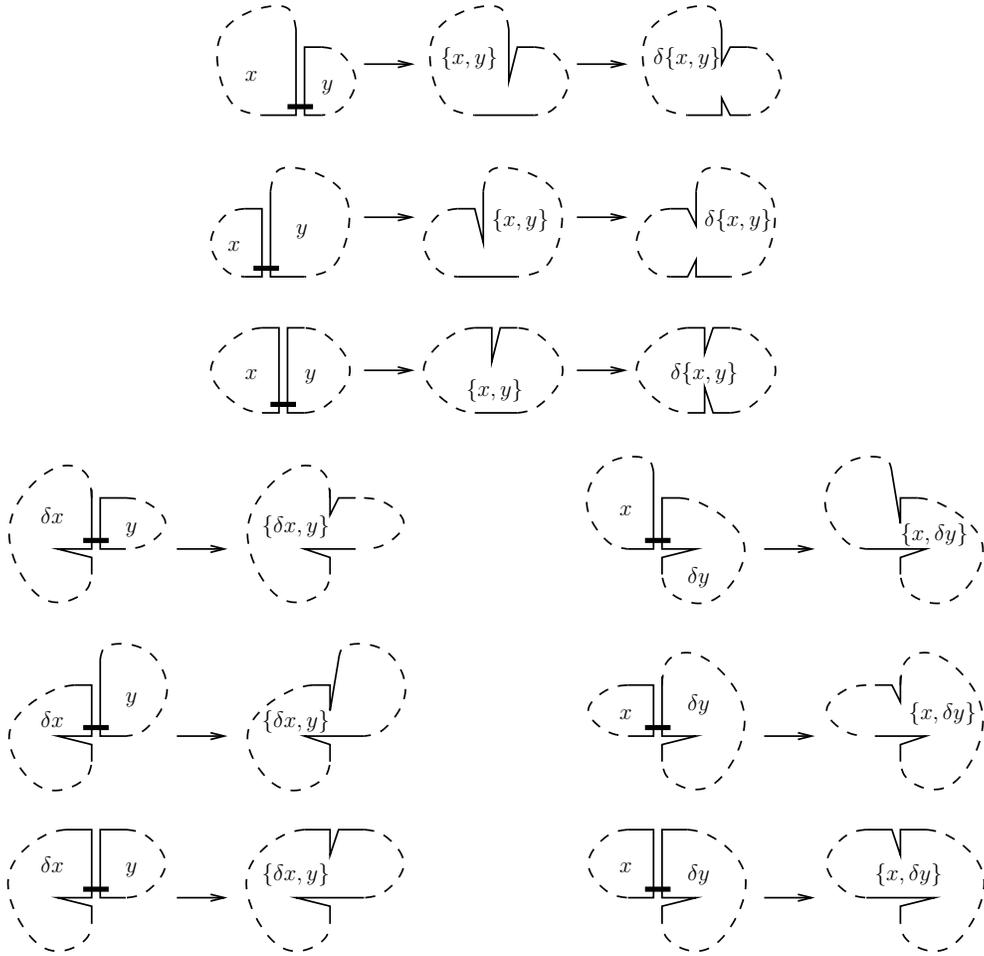


FIGURE 3.7. Exceptional terms in $f(x, y)$. The corners glued under the SFT bracket are indicated by a heavy bar. This figure only shows terms where the x corner lies counterclockwise from the y corner at the gluing; there is a corresponding set of terms where the x corner lies clockwise from the y corner, obtained by reflecting each of the pictures.

on which gluing is used. See Figure 3.9. This completes the proof of (5) mod 2.

To establish (5) over \mathbb{Z} , we just need to check signs for each of the nine pairs depicted in Figure 3.8. This is completely straightforward but somewhat tedious; we do one sample sign calculation and leave the rest to the interested reader. Suppose x, y are as depicted in Figure 3.9. Label the corners of x, y as shown in Figure 3.10. There are words w_1, w_2 such that $x = [w_1 s^*]$, the image in \mathcal{A}^{cyc} of $w_1 s^*$, and $y = s w_2 s'$. Also let ϵ_2, ϵ'_2 be the signs depicted in Figure 3.10.

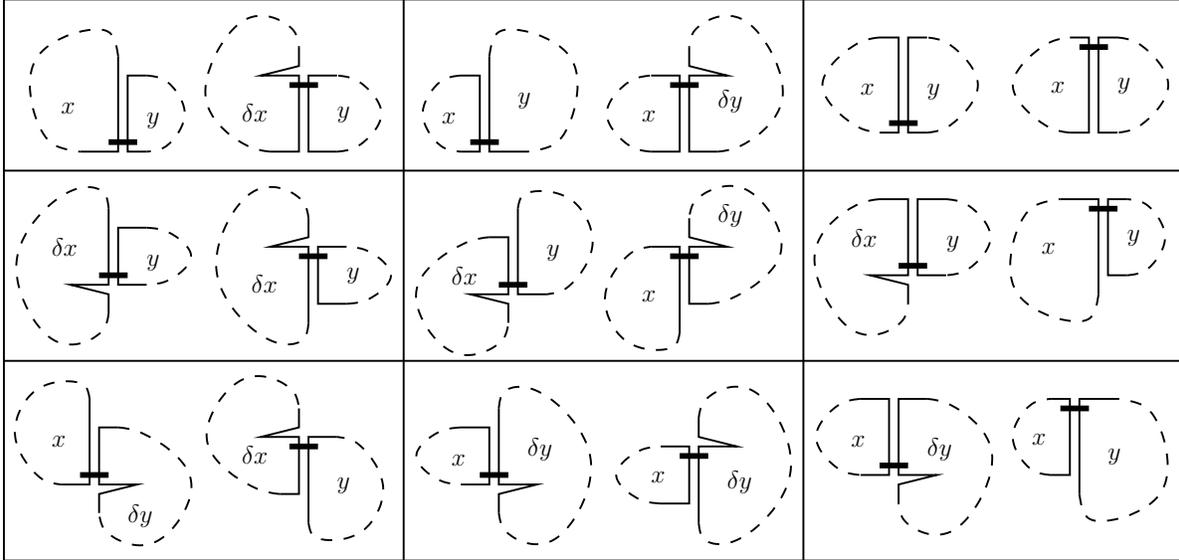


FIGURE 3.8. Pairing the exceptional terms in $f(x, y)$. Bars indicate glued corners.

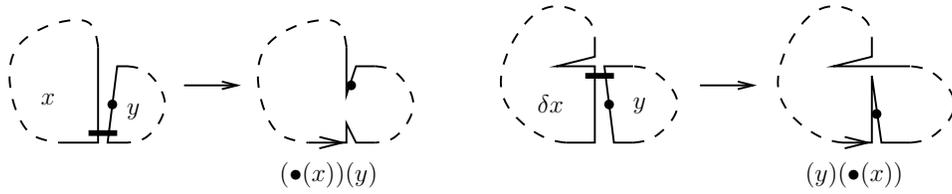


FIGURE 3.9. Two exceptional terms in $f(x, y)$ for $y \in \mathcal{A}$.

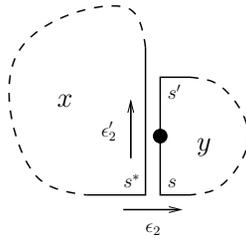


FIGURE 3.10. Labeling corners and orientations for x, y .

The relevant term in $\{x, y\}$ is $\{s^*, s\}w_1w_2s'$, where as usual $\{s^*, s\}$ is $+1$ if s is a q , -1 if s is a p . It follows that the relevant term in $\delta\{x, y\}$ is

$$\left((-1)^{|w_1|}\epsilon_2\right) w_1s^*sw_2s' = \left((-1)^{|x|}\epsilon_2'\right) xy = (\bullet(x))y,$$

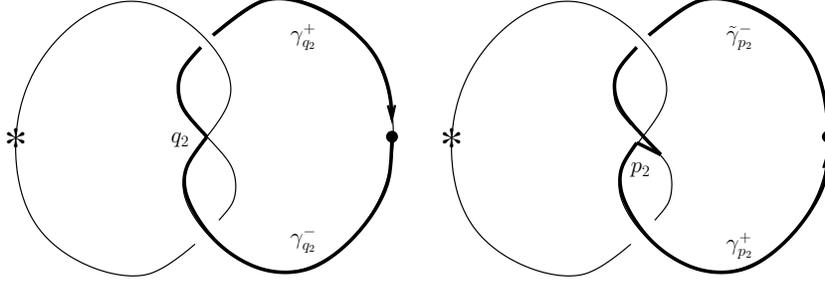


FIGURE 3.11. Based broken closed strings $\gamma_{q_2} = \gamma_{q_2}^- \cup \gamma_{q_2}^+$ (left) and $\gamma_{p_2} = \tilde{\gamma}_{p_2}^- \cup \gamma_{p_2}^+$ (right) for the knot Λ_0 .

where the first equality follows from the fact that $\epsilon_2 = (-1)^{|s^*|} \epsilon'_2$. On the other hand, the relevant term in δx is $((-1)^{|x|} \{s'^*, s'\} \epsilon'_2) w_1 s^* s'^* s'$, which is equal in \mathcal{A}^{cyc} to $((-1)^{|x|+|x||s'|} \{s'^*, s'\} \epsilon'_2) s' w_1 s^* s'^*$; thus the relevant term in $\{\delta x, y\}$ is

$$\left((-1)^{|x|+|x||s'|+|x||s w_2|} \epsilon'_2 \right) s w_2 s' w_1 s^* = (-1)^{|x||y|} y(\bullet(x)).$$

Combining the terms in $\delta\{x, y\}$ and $-\{\delta x, y\}$ contributes $[\bullet(x), y]$ to $f(x, y)$, as desired. \square

The string differential d_{str} from Section 2.3 was defined to satisfy the following result.

Proposition 3.9. *On $\hat{\mathcal{A}}$, we have $d_{\text{str}} = \delta$.*

Proof. Since δ is a derivation by Proposition 3.8(2), it suffices to show that $d_{\text{str}}(s) = \delta(s)$ when s is a generator of \mathcal{A} .

If s is a p or q , we can use the paths γ_s^\pm from Section 2.3 to define a based broken closed string γ_s with holomorphic corners such that $w(\gamma_s) = s$. More precisely, if s has holomorphic capping paths, then define $\gamma_s = \gamma_s^- \cup \gamma_s^+$; if s has antiholomorphic capping paths, then define γ_s to be a perturbation of $\gamma_s^- \cup \gamma_s^+$ to have a holomorphic corner at s (explicitly, let γ_s^- run past the crossing and return, and then join γ_s^+ to it). See Figure 3.11. If $s = t^{\pm 1}$, define paths γ_t to run along Λ once, and $\gamma_{t^{-1}}$ to run along Λ once with the reverse orientation.

In all cases, it is now trivial to check, using γ_s , that $\delta(s)$ from Definition 3.6 agrees with $d_{\text{str}}(s)$ from Definition 2.24. \square

3.3. The Hamiltonian and the LSFT algebra. Having introduced the SFT bracket and the δ map, we are now in a position to redefine the LSFT algebra $(\hat{\mathcal{A}}, d)$ in terms of strings. First we introduce the Hamiltonian $h \in \hat{\mathcal{A}}^{\text{cyc}}$ counting rigid holomorphic disks with boundary on $\Lambda \times \mathbb{R}$.

Let f be an immersed disk in \mathbb{R}^2 with boundary on $\pi_{xy}(\Lambda)$ and convex corners. More precisely, in the language of Definition 2.18, $f \in \Delta(s_1, \dots, s_k)$ for

some s_1, \dots, s_k ; recall that f is equally well an element of $\Delta(s_2, \dots, s_k, s_1)$ and other cyclic permutations as well. The boundary of f is a broken closed string in Λ with corresponding word $s_2 \cdots s_k s_1 \in \mathcal{A}$ (note that s_1 appears *last* in this word), and we can also associate a sign $\epsilon(f; s_1)$ to f , as defined in Section 2.3.

Define $\tilde{w}(f) \in \mathcal{A}^{\text{cyc}}$ to be the image in \mathcal{A}^{cyc} of $\epsilon(f; s_1) s_2 \cdots s_k s_1$. The key point is the following.

Lemma 3.10. *The element $\tilde{w}(f)$ of \mathcal{A}^{cyc} depends only on the disk f and not on which puncture is labeled s_1 ; that is, $\tilde{w}(f)$ is independent of whether f is viewed as an element of $\Delta(s_1, \dots, s_k)$, $\Delta(s_2, \dots, s_k, s_1)$, or any other cyclic permutation.*

Proof. This is clear mod 2. To check signs, suppose $f \in \Delta(s_1, \dots, s_k) = \Delta(s_2, \dots, s_k, s_1)$. If we view f as an element of $\Delta(s_1, \dots, s_k)$, then $\tilde{w}(f) = \epsilon(f; s_1) s_2 \cdots s_k s_1$, while if we view f as an element of $\Delta(s_2, \dots, s_k, s_1)$, then $\tilde{w}(f) = \epsilon(f; s_2) s_3 \cdots s_k s_1 s_2$. But we have $\epsilon(f; s_1) \epsilon(f; s_2) = \epsilon'(f; s_1) \epsilon'(f; s_2) = (-1)^{|s_2|}$, where $\epsilon'(f; s_1), \epsilon'(f; s_2)$ are the signs shown in Figure 2.5 and the second equality follows from Lemma 2.17. Since $|f| = -2$ by Lemma 2.23 (cf. Lemma 3.12 below),

$$\epsilon(f; s_1) s_2 \cdots s_k s_1 = \epsilon(f; s_1) (-1)^{|s_2|} s_3 \cdots s_k s_1 s_2 = \epsilon(f; s_2) s_3 \cdots s_k s_1 s_2$$

in \mathcal{A}^{cyc} , and the lemma follows. \square

Definition 3.11. The Hamiltonian $h \in \hat{\mathcal{A}}^{\text{cyc}}$ is the sum of $\tilde{w}(f)$ over all immersed disks f in some $\Delta(s_1, \dots, s_k)$ for all possible $k \geq 1$ and all possible s_1, \dots, s_k . (Here we mod out by cyclic permutations and count each immersed disk once; that is, we count $f \in \Delta(s_1, \dots, s_k)$ once and not k times.)

It is entirely possible that h is an infinite sum; see Figure 3.12 for an example. We do however have the following result.

Lemma 3.12. *The Hamiltonian h has degree -2 and is an element of $\mathcal{F}^1 \hat{\mathcal{A}}^{\text{cyc}}$.*

Proof. The fact that h has degree -2 can be proved in the same way as Lemma 2.23. To show that $h \in \mathcal{F}^1 \hat{\mathcal{A}}^{\text{cyc}}$, we claim that all terms in h contain at least one p , and that only finitely many terms contain at most k p 's for any k . The first part is evident from Lemma 2.22. For the second part, Lemma 2.22 implies that there are only finitely many nonempty moduli spaces of disks $\Delta(s_1, \dots, s_l)$ for which k of the s_i 's are p 's. But for fixed l and s_1, \dots, s_l , the set $\Delta(s_1, \dots, s_l)$ is finite by a standard argument given in [Che02]. \square

We now have the following result.

Proposition 3.13 (Quantum master equation). $\delta h + \frac{1}{2} \{h, h\} = 0$.

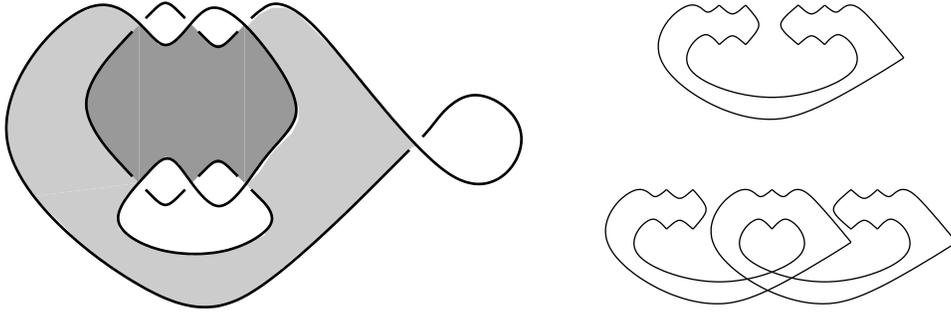


FIGURE 3.12. The xy projection of a Legendrian knot for which h is infinite. One immersed disk is indicated in the diagram on the left (the darker-shaded region is covered twice), and heuristically redrawn as the diagram on the top right (pulled apart from itself for clarity). This can be extended by wrapping more times around the bottom section of the projection, to produce an infinite family of contributions to h , the next of which is drawn on the bottom right. This knot is isotopic through restricted Reidemeister II moves to the standard Legendrian unknot.

Remark 3.14. Despite the presence of $\frac{1}{2}$ in the statement of Proposition 3.13, the result can be interpreted over \mathbb{Z} or even $\mathbb{Z}/2$. To do this, rewrite $\{x, y\}$ as a difference of two terms, $(x \rightarrow y) - (y \rightarrow x)$, where $(x \rightarrow y)$ counts terms in $\{x, y\}$ where a p in x is glued to the corresponding q in y and similarly for $(y \rightarrow x)$. Then we can write $(h \rightarrow h)$ instead of $\frac{1}{2}\{h, h\}$.

Proof of Proposition 3.13. We wish to show that $\delta h = -(h \rightarrow h)$ in the notation of Remark 3.14. First argue mod 2. As in the standard proof of

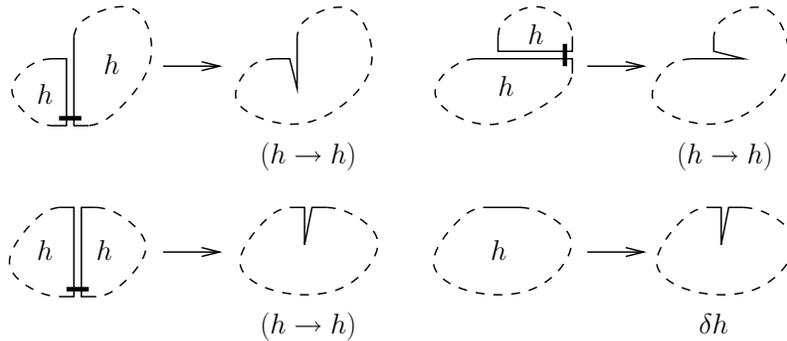


FIGURE 3.13. Paired contributions to $\delta h + (h \rightarrow h)$. Top row: an obtuse disk contributes to two canceling terms in $(h \rightarrow h)$. Bottom row: an immersed disk contributes to terms in $(h \rightarrow h)$ and δh , which also cancel.

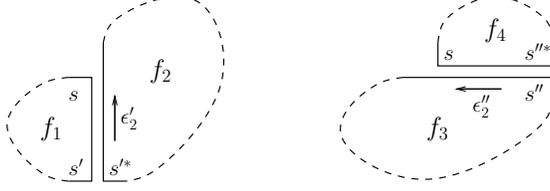


FIGURE 3.14. A closer look at obtuse disks.

$d^2 = 0$ in Chekanov [Che02], most of the terms in $(h \rightarrow h)$ cancel pairwise. Terms in $(h \rightarrow h)$ correspond to gluing two immersed disks at a corner; near this corner, the two disks overlap on an edge. If the overlapping edges are not identical, then the result is an “obtuse disk” with one concave corner, and this obtuse disk appears twice in $(h \rightarrow h)$. See the top line of Figure 3.13. If the overlapping edges are identical, then the glued disk is also an immersed disk, and the contribution of the glued disk to $(h \rightarrow h)$ is canceled by the contribution of the immersed disk to δh . See the bottom line of Figure 3.13.

As usual, to complete the proof, we need to compute signs. We claim that the two obtuse disks (the top row of Figure 3.13) give canceling contributions to $(h \rightarrow h)$; a similar calculation shows that the bottom row of Figure 3.13 gives canceling contributions to $(h \rightarrow h) + \delta h$. Consider the four disks f_1, f_2, f_3, f_4 shown in Figure 3.14. The contribution of, e.g., f_1, f_2 to $(h \rightarrow h)$ is either $\{\tilde{w}(f_1), \tilde{w}(f_2)\}$ or $\{\tilde{w}(f_2), \tilde{w}(f_1)\}$ depending on which of f_1, f_2 contains the p and which the q , but these two quantities are equal since h has degree -2 . It thus suffices to show that the contributions of $\{\tilde{w}(f_1), \tilde{w}(f_2)\}$ and $\{\tilde{w}(f_3), \tilde{w}(f_4)\}$ to $\{h, h\}$ have opposite sign.

We have

$$\tilde{w}(f_1) = (\epsilon_s(f_1)\epsilon_{s'}(f_1)\epsilon_{w_1}(f_1)\epsilon'_2) sw_1s'$$

where w_1 is some word, $\epsilon_s(f_1), \epsilon_{s'}(f_1)$ are the orientation signs for corners s, s' in f_1 , $\epsilon_{w_1}(f_1)$ is the product of orientation signs over all other corners of f_1 (i.e., the corners corresponding to w_1), and ϵ'_2 is the sign depicted in Figure 3.14 (as usual, relative to the knot orientation). Similarly, we have

$$\begin{aligned}\tilde{w}(f_2) &= (-\epsilon_{s''^*}(f_2)\epsilon_{w_2}(f_2)\epsilon'_2) s''^*w_2 \\ \tilde{w}(f_3) &= (\epsilon_{s''}(f_3)\epsilon_{w_3}(f_3)\epsilon''_2) w_3s'' \\ \tilde{w}(f_4) &= (-\epsilon_{s''^*}(f_4)\epsilon_s(f_4)\epsilon_{w_4}(f_4)) s''^*w_4s.\end{aligned}$$

Gluing s' in f_1 to s''^* in f_2 yields a contribution of

$$(-\epsilon_s(f_1)\epsilon_{s'}(f_1)\epsilon_{s''^*}(f_2)\epsilon_{w_1}(f_1)\epsilon_{w_2}(f_2)\{s', s''^*\}) sw_1w_2 = (-\epsilon_s(f_1)\epsilon_{w_1}(f_1)\epsilon_{w_2}(f_2)) sw_1w_2$$

to $\{\tilde{w}(f_1), \tilde{w}(f_2)\}$ by Lemma 2.20; similarly, gluing s'' in f_3 to s''^* in f_4 yields a contribution of $(-\epsilon_s(f_4)\epsilon_{w_3}(f_3)\epsilon_{w_4}(f_4)) w_3w_4s$ to $\{\tilde{w}(f_3), \tilde{w}(f_4)\}$. But $w_1w_2 = w_3w_4$ and hence $sw_1w_2 = w_3w_4s$ in \mathcal{A}^{cyc} , while $\epsilon_{w_1}(f_1)\epsilon_{w_2}(f_2) = \epsilon_{w_3}(f_3)\epsilon_{w_4}(f_4)$ and hence

$$(-\epsilon_s(f_1)\epsilon_{w_1}(f_1)\epsilon_{w_2}(f_2)) = -(-\epsilon_s(f_4)\epsilon_{w_3}(f_3)\epsilon_{w_4}(f_4))$$

since $\epsilon_s(f_1) = -\epsilon_s(f_4)$ by Lemma 2.20. This shows that the obtuse disks give canceling contributions to $(h \rightarrow h)$, as desired. \square

The following result, the string version of the “ $d^2 = 0$ result” Theorem 2.25, implies Theorem 2.25.

Proposition 3.15. *Define $d: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ by*

$$d(x) = \{h, x\} + \delta x.$$

Then $\{h, x\}$, δx , and $d(x)$ coincide with $d_{\text{SFT}}(x)$, $d_{\text{str}}(x)$, and $d(x)$ as defined in Section 2.3, and $(\hat{\mathcal{A}}, d)$ is an LSFT algebra with $F_d = \bullet(h)$.

Proof. We have already seen in Proposition 3.9 that $\delta = d_{\text{str}}$. The fact that $\{h, \cdot\} = d_{\text{SFT}}$ follows from a direct inspection of the definitions of h , d_{SFT} , and the SFT bracket.

It remains to show that $d^2x = [\bullet(h), x]$. But Proposition 3.4 implies that $\{h, \{h, x\}\} = \{\frac{1}{2}\{h, h\}, x\}$, whence by Propositions 3.8 and 3.13,

$$\begin{aligned} d^2x &= \{h, \{h, x\}\} + \{h, \delta x\} + \delta\{h, x\} + \delta^2x \\ &= -\{\delta h, x\} + \{h, \delta x\} + \delta\{h, x\} \\ &= [\bullet(h), x], \end{aligned}$$

as desired. \square

4. PROOF OF INVARIANCE

This section is devoted to the proof of the invariance result, Theorem 2.28. The LSFT algebra structure is associated to a Legendrian knot Λ with two marked points $*$, \bullet . Any two Legendrian-isotopic knots with marked points can be related by a sequence of four basic moves: keeping the knot fixed and sliding $*$ along it; keeping the knot fixed and sliding \bullet along it; changing the knot by a Reidemeister II move, while keeping $*$, \bullet fixed and away from the move; and changing the knot by a Reidemeister III move, while keeping $*$, \bullet fixed and away from the move.

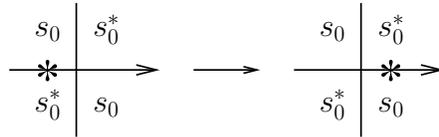


FIGURE 4.1. Changing the marked point $*$ by moving it through a crossing of $\pi_{xy}(\Lambda)$, in the direction of the orientation of Λ . There is an analogous diagram moving the marked point \bullet through the crossing.

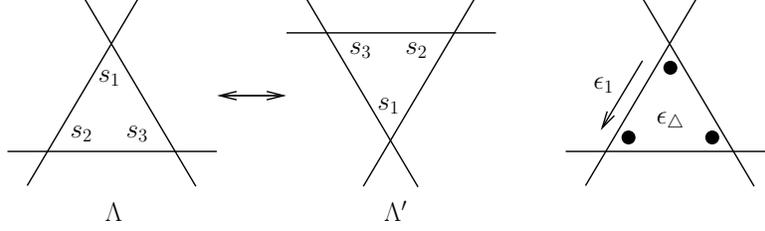


FIGURE 4.2. Reidemeister III move. The right-hand diagram defines signs $\epsilon_\Delta, \epsilon_1$: ϵ_Δ is the product of the orientation signs of the three marked corners, while ϵ_1 is the depicted orientation relative to the orientation of Λ .

Changing the marked points changes the LSFT algebra in a fairly trivial way. One can readily check from the definitions that moving $*$ across a crossing of $\pi_{xy}(\Lambda)$ labeled by s_0, s_0^* , as shown in Figure 4.1, has the effect of replacing s_0 by $t^{-1}s_0$ and s_0^* by s_0^*t , and thus corresponds to a basis change. On the other hand, moving \bullet in the same way does not change d_{SFT} but does change d_{str} by

$$d'_{\text{str}}(s) = d_{\text{str}}(s) - \{s_0, s_0^*\}[s_0 s_0^*, s]$$

for all generators s ; this corresponds to a gauge change with $z = -\{s_0, s_0^*\} s_0 s_0^*$, in the notation of Definition 2.11.

The remaining moves, Reidemeister III and Reidemeister II, are addressed in Sections 4.1 and 4.2, respectively. These are essentially extensions of the invariance arguments for the contact-homology differential graded algebra from [Che02].

4.1. Reidemeister III. Here we assume that Λ and Λ' are related by a Reidemeister III move, as shown in Figure 4.2. We may also assume without loss of generality that the points $*$, \bullet are not involved in the move and lie outside of the local pictures. Let $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}, d')$ be the LSFT algebras associated to Λ and Λ' , respectively. Note that we have identified the underlying tensor algebras by labeling the three relevant crossings as shown; all other crossings are labeled identically for the two pictures. Figure 4.2 also defines two signs $\epsilon_\Delta, \epsilon_1$.

Let ϕ be the basis change defined as follows:

$$\begin{aligned} s_1 &\mapsto s_1 + \epsilon\{s_1^*, s_1\}s_3^*s_2^* \\ s_2 &\mapsto s_2 + \epsilon\{s_2^*, s_2\}s_1^*s_3^* \\ s_3 &\mapsto s_3 + \epsilon\{s_3^*, s_3\}s_2^*s_1^* \\ t^{\pm 1} &\mapsto t^{\pm 1} \\ s &\mapsto s \end{aligned} \quad \text{for } s \neq s_1, s_2, s_3,$$

where $\epsilon = \{s_1^*, s_1\}\{s_2^*, s_2\}\{s_3^*, s_3\}\epsilon_\Delta$. Note that ϕ preserves the filtration on $\hat{\mathcal{A}}$. Indeed, the only way this would not hold would be if s_1, s_2, s_3 were all

p 's, in which case the move depicted in Figure 4.2 would not be a topological isotopy.

Lemma 4.1. *If $x \in \mathcal{A}^{\text{cyc}}$ and $y \in \mathcal{A}$, then*

$$(4) \quad \phi(\{x, y\}) = \{\phi(x), \phi(y)\}.$$

We remark that Lemma 4.1 does not hold if ϕ is an arbitrary basis change.

Proof of Lemma 4.1. By Proposition 3.4(2), it suffices to establish (4) for $y = s$, where s is any p or q . If $s \neq s_1, s_2, s_3$, then (4) holds because s does not interact with any of the quadratic terms in ϕ . It remains to prove (4) for $y = s_1, s_2, s_3$. We may assume without loss of generality that $y = s_1$ and further that x is (the cyclic quotient of) a word in \mathcal{A} .

Contributions to $\{x, y\}$ come from appearances of s_1^* in x : if $x = w_1 s_1^* w_2$, then $\{x, y\}$ contains the term $\pm w_2 w_1$, and $\phi(\{x, y\})$ contains the term $\pm \phi(w_2) \phi(w_1)$. On the other hand, contributions to $\{\phi(x), \phi(y)\} = \{\phi(x), s_1 + \epsilon\{s_1^*, s_1\} s_3^* s_2^*\}$ come from appearances of any of s_1^*, s_2, s_3 in x . The appearances of s_1^* give a contribution to $\{\phi(x), \phi(y)\}$ exactly equal to the corresponding contribution to $\phi\{x, y\}$. Any appearance of s_2 gives two canceling contributions to $\{\phi(x), \phi(y)\}$: if $x = w_1 s_2 w_2$, then $\phi(x) = \phi(w_1) s_2 \phi(w_2) + \epsilon\{s_2^*, s_2\} \phi(w_1) s_1^* s_3^* \phi(w_2)$, and the contribution to $\{\phi(x), s_1 + \epsilon\{s_1^*, s_1\} s_3^* s_2^*\}$ is

$$\begin{aligned} & (-1)^{|w_1 s_2| |w_2|} \left(\epsilon\{s_1^*, s_1\} \{s_2, s_2^*\} (-1)^{(|w_1 s_2 w_2| + 1) |s_3^*|} \right. \\ & \quad \left. + \epsilon\{s_2^*, s_2\} \{s_1^*, s_1\} (-1)^{|s_3^*| |w_2 w_1 s_1^*|} \right) s_3^* \phi(w_2) \phi(w_1) = 0, \end{aligned}$$

where the cancellation occurs since $|s_2| = |s_3^* s_1^*|$. Any appearance of s_3 in x similarly gives two canceling contributions to $\{\phi(x), \phi(y)\}$. It follows that (4) holds for $y = s_1$, as desired. \square

Lemma 4.2. *If γ is a path in Λ whose endpoints do not coincide with any of the endpoints of s_1, s_2, s_3 , then $\phi(\delta(\gamma)) = \delta(\gamma)$.*

Proof. The hypothesis of the lemma implies that s_1, s_2, s_3 only appear in $\delta(\gamma)$ in pairs, namely $\{s_1, s_1^*\} s_1^* s_1 + \{s_2, s_2^*\} s_2 s_2^*$ and its two cyclic permutations (where we cyclically permute the indices 1, 2, 3). But ϕ preserves each of these sums. \square

Lemma 4.3. *If s is a generator of $\hat{\mathcal{A}}$ not equal to $s_1, s_2, s_3, s_1^*, s_2^*, s_3^*$, then $\delta'(s) = \delta(s)$. Also,*

$$\begin{aligned} \delta'(s_1) &= \delta(s_1) - (\epsilon_1 (-1)^{|s_1|} \{s_2, s_2^*\}) s_1 s_2 s_2^* - (\epsilon_1 (-1)^{|s_1|} \{s_3^*, s_3\}) s_3^* s_3 s_1 \\ \delta'(s_1^*) &= \delta(s_1^*) + (\epsilon_1 \{s_2, s_2^*\}) s_2 s_2^* s_1^* - (\epsilon_1 \{s_3^*, s_3\}) s_1^* s_3^* s_3, \end{aligned}$$

with corresponding formulas for $\delta'(s_2), \delta'(s_3), \delta'(s_2^*), \delta'(s_3^*)$ (permute the indices cyclically).

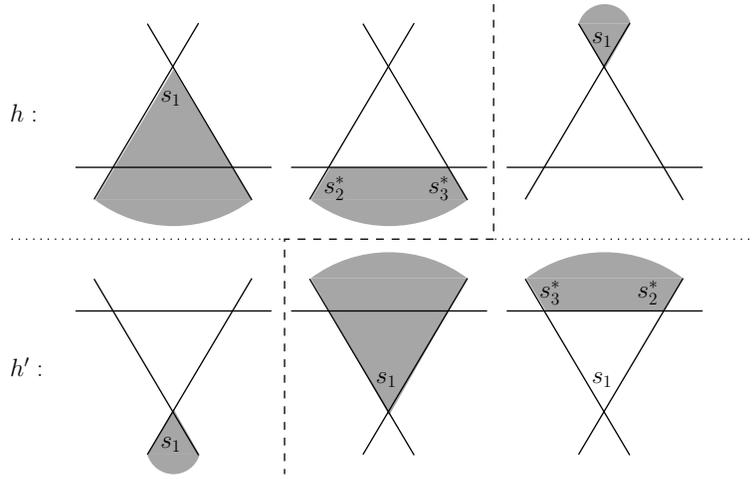


FIGURE 4.3. Contributions to h and h' for the Reidemeister III move.

Proof. The first statement is clear by Lemma 4.2. The rest follows from the definition of δ, δ' and an examination of how capping paths change under the Reidemeister III move. \square

For the next lemma, note that the triangle in the Reidemeister III move contributes the term $(\epsilon_{\Delta}\epsilon_1)s_2s_3s_1$ to both h and h' (for the latter, use Lemma 2.20).

Lemma 4.4. *Write $h_0 = h - (\epsilon_{\Delta}\epsilon_1)s_2s_3s_1$ and $h'_0 = h' - (\epsilon_{\Delta}\epsilon_1)s_2s_3s_1$; then $h'_0 = \phi(h_0)$.*

Proof. This is a standard argument along the lines of Chekanov [Che02]. Let Δ, Δ' denote the triangles bounded by s_1, s_2, s_3 in $\pi_{xy}(\Lambda), \pi_{xy}(\Lambda')$, respectively. Disks that contribute to h, h' and contain s_1 fall into two categories, depending on whether they contain Δ, Δ' or not. To a disk in h' with a corner at s_1 and not containing Δ' (the left picture in the bottom row of Figure 4.3), there are two corresponding disks in h , one with a corner at s_1 and containing Δ , the other with corners at s_3^* and s_2^* (the two left pictures in the top row). To a disk in h with a corner at s_1 and not containing Δ (right picture, top row), there are two corresponding disks in h' , one with a corner at s_1 and containing Δ' , the other with corners at s_3^* and s_2^* (two right pictures, bottom row). Similar correspondences occur for disks in h, h' with corners at s_2 or s_3 . That the signs work out follows easily from the definition of ϵ and Lemma 2.20. \square

Proposition 4.5. *Under a Reidemeister III move, the LSFT algebra changes by a change of basis. More precisely, $d' = \phi \circ d \circ \phi^{-1}$.*

Proof. It suffices to show that

$$(5) \quad \phi d(s) = d' \phi(s)$$

for s any generator of $\hat{\mathcal{A}}$.

CASE 1: $s \neq s_1, s_2, s_3, s_1^*, s_2^*, s_3^*$. By Lemma 4.3, $\phi(\delta(s)) = \delta(s) = \delta'(s)$, while by Lemmas 4.1 and 4.4, $\phi\{h_0, s\} = \{h'_0, s\}$. It follows that

$$\phi d(s) = \phi\{h_0 + (\epsilon_\Delta \epsilon_1) s_2 s_3 s_1, s\} + \phi \delta(s) = \phi\{h_0, s\} + \delta'(s) = d' \phi(s),$$

and (5) holds.

CASE 2: $s = s_1^*, s_2^*, s_3^*$. By symmetry, we may assume that $s = s_1^*$. Then

$$\begin{aligned} \phi d s_1^* &= \phi\{h_0, s_1^*\} + \epsilon_\Delta \epsilon_1 \phi\{s_2 s_3 s_1, s_1^*\} + \phi \delta(s_1^*) \\ &= \{h'_0, s_1^*\} + \epsilon_\Delta \epsilon_1 \{s_1, s_1^*\} (s_2 + \epsilon\{s_2^*, s_2\} s_1^* s_3^*) (s_3 + \epsilon\{s_3^*, s_3\} s_2^* s_1^*) + \phi \delta(s_1^*) \end{aligned}$$

and

$$\begin{aligned} d' \phi s_1^* &= d' s_1^* = \{h'_0, s_1^*\} + \epsilon_\Delta \epsilon_1 \{s_2 s_3 s_1, s_1^*\} + \delta'(s_1^*) \\ &= \{h'_0, s_1^*\} + \epsilon_\Delta \epsilon_1 \{s_1, s_1^*\} s_2 s_3 + \delta(s_1^*) + \epsilon_1 \{s_2, s_2^*\} s_2 s_2^* s_1^* - \epsilon_1 \{s_3^*, s_3\} s_1^* s_3^* s_3; \end{aligned}$$

thus to establish (5) for $s = s_1^*$, it suffices to show that

$$\phi \delta(s_1^*) - \delta(s_1^*) = \epsilon_1 \epsilon s_1^* s_3^* s_2^* s_1^*.$$

Because of the form of ϕ , one can disregard all terms in $\delta(s_1^*)$ except those of the form $s_1^* s_1 s_1^*$, $s_1^* s_3^* s_3$, $s_2 s_2^* s_1^*$. It is straightforward to check that the total contribution of these terms to $\phi \delta(s_1^*) - \delta(s_1^*)$ is precisely $\epsilon_1 \epsilon s_1^* s_3^* s_2^* s_1^*$ for any of the four possible configurations of $\gamma_{s_1^*}$ near s_1^* .

CASE 3: $s = s_1, s_2, s_3$. We may assume that $s = s_1$. Now

$$\phi d(s_1) = \phi\{h, s_1\} + \phi \delta(s_1) = \phi\{h_0, s_1\} + \phi \delta(s_1)$$

and

$$\begin{aligned} d' \phi(s_1) &= \{h'_0, \phi(s_1)\} + \epsilon_\Delta \epsilon_1 \{s_2 s_3 s_1, s_1 + \epsilon\{s_1^*, s_1\} s_3^* s_2^*\} + \delta'(s_1 + \epsilon\{s_1^*, s_1\} s_3^* s_2^*) \\ &= \phi\{h_0, s_1\} + \delta'(s_1 + \epsilon\{s_1^*, s_1\} s_3^* s_2^*) \\ &\quad + \epsilon_1 \{s_2^*, s_2\} \{s_3^*, s_3\} \left((-1)^{|s_1|} \{s_3, s_3^*\} s_1 s_2 s_2^* + (-1)^{|s_2|+|s_3^*|} \{s_2, s_2^*\} s_3^* s_3 s_1 \right) \\ &= \phi\{h_0, s_1\} + \delta(s_1) + \epsilon\{s_1^*, s_1\} \delta'(s_3^* s_2^*) \end{aligned}$$

by Lemmas 4.1, 4.3, and 4.4; thus to establish (5) for $s = s_1$, it suffices to show that

$$(6) \quad \phi \delta(s_1) - \delta(s_1) = \epsilon\{s_1^*, s_1\} \delta'(s_3^* s_2^*).$$

By Lemma 4.2, (6) simply states that replacing each appearance of s_1 in $\delta(s_1)$ by $s_3^* s_2^*$ results in $\delta'(s_3^* s_2^*)$. But given a based broken closed string in Λ with a single holomorphic corner at s_1 , a small perturbation yields a based broken closed string in Λ' whose word is $s_3^* s_2^*$; the correspondence between these strings yields (6). \square

4.2. Reidemeister II. Here we assume that Λ and Λ' are related by a Reidemeister II move, as shown in Figure 4.4. At some point it will become important that the move is a restricted Reidemeister II move; we will indicate where we use this fact in the proof.

As in the Reidemeister III case, assume that the points $*$, \bullet are not involved in the move and lie outside of the local pictures. Let $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}', d')$ be the LSFT algebras associated to Λ and Λ' , respectively. View the algebra $\hat{\mathcal{A}}$ as a stabilization of $\hat{\mathcal{A}}'$ by adding four generators p_a, q_a, p_b, q_b corresponding to the two new crossings in Λ . Then we can extend d' to $\hat{\mathcal{A}}$ by setting

$$\begin{aligned} d'(q_a) &= q_b \\ d'(q_b) &= [F_{d'}, q_a] \\ d'(p_b) &= p_a \\ d'(p_a) &= [F_{d'}, p_b]. \end{aligned}$$

Note that this makes $(\hat{\mathcal{A}}, d')$ an LSFT algebra and a stabilization of $(\hat{\mathcal{A}}', d')$.

We claim that $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}, d')$ are related by a basis change; this will prove invariance under restricted Reidemeister II.

The bigon in Figure 4.4 contributes the term $\epsilon p_a q_b$ to the Hamiltonian for Λ , where ϵ is the sign depicted in Figure 4.4; this then contributes $-\epsilon p_a, \epsilon(-1)^{|p_a|} q_b$ to $d(p_b), d(q_a)$ respectively. Write

$$\begin{aligned} d(p_b) &= -\epsilon p_a + u \\ d(q_a) &= \epsilon(-1)^{|p_a|} q_b + v \end{aligned}$$

for some $u, v \in \hat{\mathcal{A}}$. Let ϕ_0 be the algebra map on $\hat{\mathcal{A}}$ defined by

$$\begin{aligned} \phi_0(q_b) &= \epsilon(-1)^{|p_a|} q_b + v \\ \phi_0(p_a) &= -\epsilon p_a + u \\ \phi_0(s) &= s \quad \text{for } s \neq q_b, p_a. \end{aligned}$$

Lemma 4.6. ϕ_0 is a basis change on $\hat{\mathcal{A}}$.

Proof. We have $\phi_0 = \phi_b \circ \phi_a$, where ϕ_b is the elementary automorphism supported at q_b that sends q_b to $\epsilon(-1)^{|p_a|} q_b + v$, and ϕ_a is the elementary

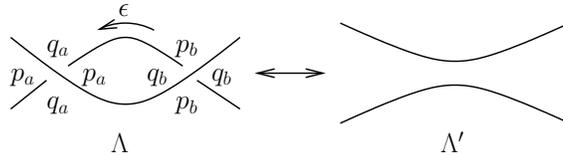


FIGURE 4.4. Reidemeister II move. The two crossings in $\pi_{xy}(\Lambda)$ are given labels p_a, q_a, p_b, q_b in their quadrants, as shown. A sign ϵ is also shown.

automorphism supported at p_a that sends p_a to $-\epsilon p_a + \phi_b^{-1}(u)$. Now by Lemma 2.22, any term in v either involves a p or only involves $t^{\pm 1}$ and q 's of smaller height than q_b , and so ϕ_b is an elementary automorphism of $\hat{\mathcal{A}}$. Also by Lemma 2.22, any term in u either involves two p 's or only involves $t^{\pm 1}$, q 's, and p 's of greater height than p_a ; it follows that the only terms in $\phi_b^{-1}(u)$ not containing at least two p 's do not involve p_a , and so ϕ_a is an elementary automorphism as well. \square

So far, the argument given here is a straightforward extension of Chekanov's proof of Reidemeister-II invariance in [Che02, section 8.4], which then hinges on the following two points [Che02, Lemma 8.2]:

- (1) d' and $\phi_0^{-1} \circ d \circ \phi_0$ agree on q generators whose height is at most the height of q_b ;
- (2) on all q generators, d' and $\phi_0^{-1} \circ d \circ \phi_0$ agree modulo terms involving either q_a or q_b .

Chekanov uses these two properties to bootstrap ϕ_0 up to an automorphism that intertwines d' and d for all q generators.

The proof in the current circumstance is complicated by the fact that (1) no longer holds, due to the possible presence of p 's in the differentials of any q generator. Nevertheless, an analogue of (2) still holds and is presented as Lemma 4.9 below. We will use this, along with a property of the differential that we call "ordered"-ness (Definition 4.10) that loosely generalizes property (1), to perform a bootstrapping argument similar to Chekanov's.

In order to prove the analogue of (2), we need to establish a lemma that extends the central argument in Chekanov's proof of (2) [Che02, section 8.5]. Define a graded algebra map $\psi: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ by $\psi(p_a) = \epsilon u$, $\psi(q_b) = -\epsilon(-1)^{|p_a|}v$, $\psi(q_a) = \psi(p_b) = 0$, and $\psi(s) = s$ for all other generators of $\hat{\mathcal{A}}$.

Lemma 4.7. *Let s be a generator of $\hat{\mathcal{A}}$ besides $p_a, q_a, p_b, q_b, t^{\pm 1}$, and let ψ^n denote the n -th iterate of ψ . In $\hat{\mathcal{A}}$, the limit $\lim_{n \rightarrow \infty} \psi^n d_{\text{SFT}}(s)$ exists and equals $d'_{\text{SFT}}(s)$.*

Proof. We first show that the limit exists. We can assume that $h(p_a) - h(p_b) > 0$ is arbitrarily small, since by Lemma 2.22 this is the area of the bigon determined by p_a and q_b . By Lemma 2.22 again, this implies that any term in v or u involving q_b must be $O(p)$, while any term in u involving p_a must be $O(p^2)$. It follows that for any m , the portion of $\psi^n d_{\text{SFT}}(s)$ involving at most m p 's stabilizes as $n \rightarrow \infty$, and thus the limit exists.

It remains to show that

$$(7) \quad \lim_{n \rightarrow \infty} \psi^n d_{\text{SFT}}(s) = d'_{\text{SFT}}(s).$$

First assume for simplicity that $u, v \in \hat{\mathcal{A}}'$, i.e., that u, v do not involve p_a, q_a, p_b, q_b . In this case, the left hand side of (7) is $\psi d_{\text{SFT}}(s)$.

Any term in $d_{\text{SFT}}(s)$ not involving any of p_a, q_a, p_b, q_b corresponds to a disk preserved by the Reidemeister II move and thus appears in $d'_{\text{SFT}}(s)$

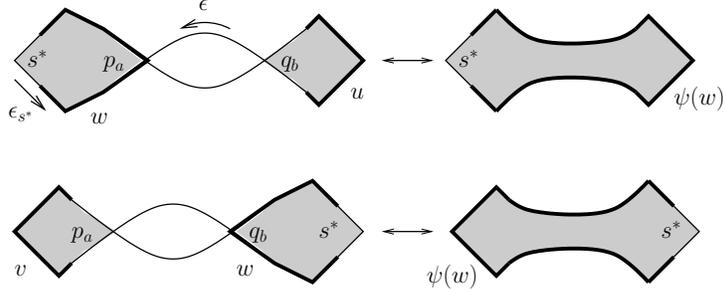


FIGURE 4.5. Top: gluing a term w in $d_{\text{SFT}}(s)$ to u to obtain a term (or sum of terms) $\psi(w)$ in $d'_{\text{SFT}}(s)$; bottom: gluing a term w in $d_{\text{SFT}}(s)$ to v to obtain a term $\psi(w)$ in $d'_{\text{SFT}}(s)$. Two signs ϵ, ϵ_{s^*} are also shown.

as well. Any term in $d_{\text{SFT}}(s)$ involving either q_a or p_b is killed by ψ . The remaining terms in $d_{\text{SFT}}(s)$ involve p_a or q_b but not q_a or p_b ; call such a term w . Then $\psi(w)$ appears in $d'_{\text{SFT}}(s)$: at each p_a or q_b corner of w , glue u or v , respectively; this gives the disks in $d'_{\text{SFT}}(s)$ passing through the neck in $\pi_{xy}(\Lambda')$. See Figure 4.5. This proves (7) mod 2 in this case.

In fact, the signs in the definition of ψ work out so that $d'_{\text{SFT}}(s) = \psi d_{\text{SFT}}(s)$ over \mathbb{Z} . Consider for instance gluing u to p_a , where we abuse notation and use u to denote a particular term in u . Let ϵ_w, ϵ_u be the product of the orientation signs over the corners in w, u , and let ϵ_{s^*} be the sign shown in Figure 4.5. Then w appears in $d_{\text{SFT}}(s)$ with sign $\epsilon_{s^*}\epsilon_w$, while u has sign $\epsilon\epsilon_u$. On the other hand, $\psi(w)$ appears in $d'_{\text{SFT}}(s)$ with sign $\epsilon_{s^*}\epsilon_w\epsilon_u$ since the orientation signs for the corners of w, u at p_a, q_b , respectively, are equal. This agrees with the fact that ψ sends p_a to ϵu . We can similarly check the sign in $\psi(q_b) = -\epsilon(-1)^{|p_a|}v$. This completes the proof of (7) when $u, v \in \hat{\mathcal{A}}'$.

In general, even if u, v involve p_a, q_a, p_b, q_b , the same argument shows (7). To get $d'_{\text{SFT}}(s)$, one starts with $d_{\text{SFT}}(s)$ and keeps replacing any appearances of p_a, q_b, q_a, p_b by $\epsilon u, -\epsilon(-1)^{|p_a|}v, 0, 0$, respectively; but this is precisely what $\lim_{n \rightarrow \infty} \psi^n d_{\text{SFT}}(s)$ gives. \square

Definition 4.8. If $x, y \in \hat{\mathcal{A}}$, write $x \equiv y \pmod{a, b}$ if $x - y$ only includes terms involving at least one of p_a, q_a, p_b, q_b (in the notation from Section 2.2, $\pi(x - y) = 0$). If f, g are two maps from $\hat{\mathcal{A}}$ to $\hat{\mathcal{A}}$, write $f \equiv g \pmod{a, b}$ if $f(x) \equiv g(x) \pmod{a, b}$ for all $x \in \hat{\mathcal{A}}$.

Lemma 4.9. On $\hat{\mathcal{A}}$, $d' \equiv \phi_0^{-1} \circ d \circ \phi_0 \pmod{a, b}$.

Proof. It suffices to show that $d'(s) \equiv \phi_0^{-1} \circ d \circ \phi_0(s) \pmod{a, b}$ for all generators s of $\hat{\mathcal{A}}$. We have

$$\phi_0^{-1} d \phi_0(p_a) = \phi_0^{-1} d^2(p_b) = \phi_0^{-1} [F_d, p_b] \equiv 0 \equiv [F_{d'}, p_b] = d'(p_a) \pmod{a, b}$$

and similarly $\phi_0^{-1}d\phi_0(s) \equiv d'(s) \equiv 0 \pmod{a, b}$ for $s = q_a, p_b, q_b$. Also, $d'(t) - d(t)$ consists of four terms, one each involving $p_a q_a, q_a p_a, p_b q_b, q_b p_b$, and it follows easily that $\phi_0^{-1}d\phi_0(t^{\pm 1}) \equiv d'(t^{\pm 1}) \pmod{a, b}$.

Now assume that $s \neq p_a, q_a, p_b, q_b, t^{\pm 1}$; we want to show that $d's \equiv \phi_0^{-1}ds \pmod{a, b}$. Since $d'_{\text{str}}s \equiv \phi_0^{-1}d_{\text{str}}s \pmod{a, b}$ as before, it suffices to show that $d'_{\text{SFT}}(s) \equiv \phi_0^{-1}d_{\text{SFT}}(s) \pmod{a, b}$. By Lemma 4.7, this follows from establishing that $\lim_{n \rightarrow \infty} \psi^n \equiv \phi_0^{-1} \pmod{a, b}$.

In fact, we claim that on $\hat{\mathcal{A}}$, $\lim_{n \rightarrow \infty} \psi^n = \pi \circ \phi_0^{-1}$, or equivalently $\lim_{n \rightarrow \infty} \psi^n \circ \phi_0 = \pi$, where π is the projection map from $\hat{\mathcal{A}}$ to $\hat{\mathcal{A}}'$ as usual. Indeed, since both sides are algebra maps, it suffices to check that $\lim_{n \rightarrow \infty} \psi^n \circ \phi_0(s) = \pi(s)$ for all generators s of $\hat{\mathcal{A}}$. This holds trivially unless $s = q_b$ or p_a ; it also holds for $s = q_b$ since

$$\epsilon(-1)^{|p_a|} \lim_{n \rightarrow \infty} \psi^n \phi_0(q_b) = \lim_{n \rightarrow \infty} \psi^n (q_b + \epsilon(-1)^{|p_a|} v) = \lim_{n \rightarrow \infty} (\psi^n q_b - \psi^{n+1} q_b) = 0,$$

and similarly for $s = p_a$. \square

Write $d_0 = \phi_0^{-1} \circ d \circ \phi_0$ on $\hat{\mathcal{A}}$. We claim that d_0 is related to d' by a basis change; this will imply that d is related to d' by a basis change, which will complete the invariance proof for the restricted Reidemeister II move.

Definition 4.10. We say that a derivation d on $\hat{\mathcal{A}}$ is *ordered* if

$$\begin{aligned} d(q_j) &= (\text{function of } t^{\pm 1}, q_a, q_b, q_1, \dots, q_{j-1}) + O(p) \\ d(p_j) &= (\text{function of } t^{\pm 1}, p_a, q_a, p_b, q_b, q_1, \dots, q_n, p_{j+1}, \dots, p_n) + O(p^2) \\ d(t) &= O(p). \end{aligned}$$

Order the crossings of Λ (or Λ') in such a way that $h(q_1) \leq h(q_2) \leq \dots \leq h(q_n)$; then by Lemma 2.22, d is automatically ordered.

Lemma 4.11. $d_0 = \phi_0^{-1} \circ d \circ \phi_0$ is ordered.

Proof. Since ϕ_0 preserves the p filtration, it is clear that $d_0(t) = O(p)$. Next consider $d_0(q_j) = \phi_0^{-1}(d(q_j))$, and note that ϕ_0^{-1} fixes all generators of $\hat{\mathcal{A}}$ except for q_b and p_a . We wish to show that the order p^0 terms (that is, the terms that are not $O(p^1)$) in $d_0(q_j)$ do not involve q_j, \dots, q_n . Since d is ordered and ϕ_0 fixes all words that do not involve q_b or p_a , it suffices to show that if w is a word in $d(q_j)$ involving q_b , then the order p^0 terms in $\phi_0^{-1}(w)$ do not involve q_j, \dots, q_n .

We may assume that $|h(q_j)| > |h(q_b)|$, since otherwise any term in $d(q_j)$ involving q_b must be $O(p)$ by Lemma 2.22. As in the proof of Lemma 4.9, we may also assume that $|h(q_a)| - |h(q_b)| > 0$ is arbitrarily small (more precisely, that no $|h(q_j)|$ lies in the interval $[|h(q_b)|, |h(q_a)|]$). Then since $d(q_a) = q_b + v$, any term in v involving q_j, \dots, q_n must be $O(p)$; otherwise,

by Lemma 4.9, $|h(q_a)| > |h(q_i)|$ for some $i \geq j$, and $|h(q_i)| \geq |h(q_j)|$. But then to order p^0 , v and hence $\phi_0^{-1}(w)$ does not involve q_j, \dots, q_n .

Finally, we need to prove that the order p^1 terms (that is, the terms that are not $O(p^2)$) in $d_0(p_j) = \phi_0^{-1}(d(p_j))$ do not involve p_1, \dots, p_{j-1} . As before, since d is ordered, the only place p_1, \dots, p_{j-1} can appear in the order p^1 terms in $d_0(p_j)$ is in contributions from $\phi_0^{-1}(q_b)$ or $\phi_0^{-1}(p_a)$. Any contribution from $\phi_0^{-1}(q_b)$ is $O(p^2)$, since it contains one of p_1, \dots, p_{j-1} along with some other p (from the fact that $d(p_j) = O(p)$). Now if $h(p_j) > h(p_a)$, then all terms in $d(p_j)$ involving p_a are $O(p^2)$, while if $h(p_j) < h(p_a)$, then any term in u involving p_1, \dots, p_{j-1} is $O(p^2)$ by Lemma 2.22 again. In either case, we conclude that the order p^1 terms in $\phi_0^{-1}(d(p_j))$ cannot involve p_1, \dots, p_{j-1} . \square

To find the basis change relating d_0 to d' , we need to use the fact that the Reidemeister II move is restricted. Let p_ℓ, q_ℓ label the crossing whose loop lies outside the rest of the diagram and does not interact with the move, and choose \bullet to lie on this loop and $*$ not to lie on this loop.

Lemma 4.12. *For a restricted Reidemeister II move, we have:*

- $d(p_\ell) = d'(p_\ell) = 0$;
- $F_d = F_{d'} = \epsilon_\ell p_\ell$, where ϵ_ℓ is the sign shown in Figure 4.6;
- $d'(s) = d_0(s)$ if $s = p_a, q_a, p_b$, or q_b .

Proof. The only term in the Hamiltonians for Λ or Λ' involving q_ℓ is p_ℓ . The first two properties then follow from the definitions of d, d' and Proposition 3.15. The third property can be proven by trivially modifying the first paragraph of the proof of Lemma 4.9, where we now use the fact that $\phi_0^{-1}(F_d) = F_d = F_{d'}$. \square

For ease of notation, define q_{-n}, \dots, q_0 by $q_{-j} = p_j$ for $1 \leq j \leq n$, $q_0 = t$. Then the condition for a derivation d to be ordered is that for all j with $-n \leq j \leq n$, any term in $d(q_j)$ involving one of q_j, q_{j+1}, \dots, q_n also involves

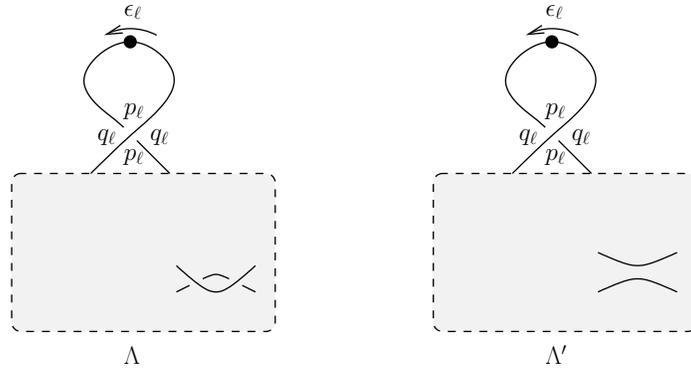


FIGURE 4.6. Labels for a restricted Reidemeister II move.

another p . Here a word w involving q_i involves another p if $i \geq 0$ and $w = O(p)$, or if $i < 0$ and $w = O(p^2)$.

Starting with $d_{0,1} := d_0$, we will inductively define a sequence of differentials $d_{j,1}, \phi_{j,1}$ for $1 \leq j \leq n$ and $d_{j,k}, \phi_{j,k}$ for $-n \leq j \leq n$ and $k \geq 2$.

Claim. We can inductively construct $d_{j,k}, \phi_{j,k}$ to satisfy the following properties:

(i) we have

$$\begin{aligned} d_{j,1} &= \phi_{j,1}^{-1} \circ d_{j-1,1} \circ \phi_{j,1} & 1 \leq j \leq n \\ d_{j,k} &= \phi_{j,k}^{-1} \circ d_{j-1,k} \circ \phi_{j,k} & 1-n \leq j \leq n, k \geq 2 \\ d_{-n,k} &= \phi_{-n,k}^{-1} \circ d_{n,k-1} \circ \phi_{-n,k} & k \geq 2; \end{aligned}$$

(ii) $d_{j,k}$ is ordered;

(iii) $d_{j,k}(s) = d'(s)$ for $s = p_a, q_a, p_b, q_b$, $d_{j,k}(p_\ell) = 0$, and $F_{d_{j,k}} = \epsilon_\ell p_\ell$;

(iv) $\phi_{j,k} \equiv \text{Id} \pmod{p^{k-1}}$ for all j, k ;

(v) $d_{j,k} \equiv d' \pmod{a, b}$ and $d_{j,k} \equiv d' \pmod{p^{k-1}}$ for all j, k ;

(vi) $d_{j,k}(q_i) \equiv d'(q_i) \pmod{p^k}$ for $i \leq j$, for all j, k .

Note that $d_{0,1}$ satisfies (ii) by Lemma 4.11, (iii) by Lemma 4.12, (v) by Lemma 4.9, and (vi) because $d_{0,1}(q_i) \equiv d'(q_i) \equiv 0 \pmod{p^1}$ for $i \leq 0$.

The following diagram summarizes the inductive order of the construction. Given $d_{j-1,k}$ for $j \geq 1-n$, we construct $\phi_{j,k}, d_{j,k}$; given $d_{n,k-1}$, we construct $\phi_{-n,k}, d_{-n,k}$. Each differential agrees with d' to the specified order in p when applied to a generator corresponding to its column or any column to its left, and to order one less in p when applied to any generator corresponding to a column to its right.

	p_n	p_{n-1}	\cdots	p_1	t	q_1	q_2	\cdots	q_n
$O(p^1)$	$d_0 = d_{0,1} \rightarrow d_{1,1} \rightarrow d_{2,1} \rightarrow \cdots \rightarrow d_{n,1}$								
$O(p^2)$	$\hookrightarrow d_{-n,2} \rightarrow d_{1-n,2} \rightarrow \cdots \rightarrow d_{-1,2} \longrightarrow d_{0,2} \longrightarrow d_{1,2} \rightarrow d_{2,2} \rightarrow \cdots \rightarrow d_{n,2}$								
$O(p^3)$	$\hookrightarrow d_{-n,3} \longrightarrow \cdots$								

Proof of Claim. Suppose that $d_{j-1,k}$ satisfies (ii) through (vi) for some $j \leq n$. Define the elementary automorphism $\phi_{j,k}$ on $\hat{\mathcal{A}}$, supported at q_j , by

$$\phi_{j,k}(q_j) = q_j - H d_{j-1,k}(q_j),$$

where $H: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ is the operator defined in Section 2.2. (Note that the notation has changed slightly from Section 2.2: $\hat{\mathcal{A}}$ there is $\hat{\mathcal{A}}'$ here, and $S_i \hat{\mathcal{A}}$ there is $\hat{\mathcal{A}}$ here.) Observe that $\phi_{j,k}$ is elementary because $d_{j-1,k}$ is ordered by assumption.

Define $d_{j,k} = \phi_{j,k}^{-1} \circ d_{j-1,k} \circ \phi_{j,k}$. We wish to show that $d_{j,k}, \phi_{j,k}$ satisfy (ii) through (vi). (A corresponding construction produces $d_{-n,k}$ from $d_{n,k-1}$; here $\phi_{-n,k}$ is supported at $q_{-n} = p_n$ and $\phi_{-n,k}(q_{-n}) = q_{-n} - Hd_{n,k-1}(q_{-n})$. The proof that $d_{-n,k}, \phi_{-n,k}$ satisfy (ii) through (vi) is entirely similar and will be omitted here.)

We first check (ii): for all i , any term in $d_{j,k}(q_i)$ involving q_i, \dots, q_n must include another p as well. If $i < j$, then since $d_{j-1,k}$ is ordered, any term in $d_{j-1,k}(q_i)$ involving q_j must include another p , and the condition holds. If $i \geq j$, then since any term in $Hd_{j-1,k}(q_j)$ involving q_j, \dots, q_n must include another p , it follows that any term in $\phi_{j,k}^{-1}(q_j)$ involving q_j, \dots, q_n must also include another p , and the condition holds here as well. This demonstrates (ii) for $d_{j,k}$.

As for the other conditions, note that (iii) holds for $d_{j,k}$ since $\phi_{j,k}$ preserves $p_a, q_a, p_b, q_b, p_\ell$ by construction (for p_ℓ , use the induction hypothesis $d_{j-1,k}(p_\ell) = 0$). Since H preserves the p filtration and $d_{j-1,k}(q_j) \equiv d'(q_j) \pmod{p^{k-1}}$, we have

$$Hd_{j-1,k}(q_j) \equiv Hd'(q_j) = 0 \pmod{p^{k-1}}$$

and thus $\phi_{j,k} \equiv \text{Id} \pmod{p^{k-1}}$, as required by (iv). It follows that (v) holds for $d_{j,k}$ since it holds for $d_{j-1,k}$. As for (vi), if $i < j$, then

$$d_{j,k}(q_i) = \phi_{j,k}^{-1} d_{j-1,k} \phi_{j,k}(q_i) = \phi_{j,k}^{-1} d_{j-1,k}(q_i) \equiv \phi_{j,k}^{-1} d'(q_i) \equiv d'(q_i) \pmod{p^k},$$

where the second-to-last equality holds since $d_{j-1,k}(q_i) \equiv d'(q_i) \pmod{p^k}$ by induction assumption, and the final equality holds because the only terms in $d'(q_i)$ that can involve q_j must also involve another p .

Thus to complete the induction step, we need to establish that $d_{j,k}(q_j) \equiv d'(q_j) \pmod{p^k}$. Since the only terms $d'(q_j)$ involving q_j itself must also involve another p , we have $\phi_{j,k}^{-1} d'(q_j) \equiv d'(q_j) \pmod{p^k}$. By the construction of $d_{j,k}$ in terms of $d_{j,k-1}$, it now suffices to show that

$$(8) \quad d_{j-1,k} \phi_{j,k}(q_j) \equiv d'(q_j) \pmod{p^k}.$$

For ease of notation, write $\tilde{d} = d_{j-1,k}$. Recall the map $H: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ from Lemma 2.16; by Lemma 2.16, we have $Hd' + d'H = \text{Id} - \iota \circ \pi$, where $\iota \circ \pi$ is the map on $\hat{\mathcal{A}}$ that projects away any term involving p_a, q_a, p_b, q_b . It follows that

$$Hd'\tilde{d}q_j + d'H\tilde{d}q_j = \tilde{d}q_j - \iota\pi\tilde{d}q_j = \tilde{d}q_j - d'q_j$$

where the last equality holds by property (v). Thus

$$\begin{aligned} \tilde{d}\phi_{j,k}q_j &= \tilde{d}(q_j - H\tilde{d}q_j) \\ &= d'q_j - (\tilde{d} - d')H\tilde{d}q_j - H(\tilde{d} - d')\tilde{d}q_j + H[F_{\tilde{d}}, q_j]. \end{aligned}$$

Now from properties (ii), (v), and (vi) for \tilde{d} , it follows that

$$(\tilde{d} - d')\tilde{d}q_j \equiv (\tilde{d} - d')H\tilde{d}q_j \equiv 0 \pmod{p^k}.$$

Also $F_{\tilde{d}} = F_{d'} \in \hat{\mathcal{A}}'$ so $H[F_{\tilde{d}}, q_j] = 0$. The desired equation (8) now follows.

This completes the induction step and the proof of the claim. \square

To finish the proof of invariance under restricted Reidemeister II, we note that since $d_0 = \phi_0^{-1} \circ d \circ \phi_0$, we can write $d' = \phi^{-1} \circ d \circ \phi$, where

$$\phi = \phi_0 \phi_{1,1} \phi_{2,1} \cdots \phi_{n,1} \phi_{-n,2} \cdots \phi_{n,2} \phi_{-n,3} \cdots .$$

This is an infinite composition, but for any k , all but finitely many terms in this composition are congruent to the identity $(\text{mod } p^k)$. For ϕ to be a change of basis, we need rewrite this as a composition of finitely many elementary automorphisms. The following result thus completes the invariance proof.

Lemma 4.13. *If, for $-n \leq j \leq n$ and $k \geq 2$, $\phi_{j,k}$ are elementary automorphisms with $\phi_{j,k}$ supported at q_j and $\phi_{j,k} \equiv \text{Id} \pmod{p^{k-1}}$ for all j, k , then we can write the infinite composition*

$$\phi_{-n,2} \cdots \phi_{n,2} \phi_{-n,3} \cdots$$

as a finite composition $\phi_{(-n)} \cdots \phi_{(n)}$, where $\phi_{(j)}$ is an elementary automorphism supported at q_j .

Proof. Consider two elementary automorphisms ϕ_1, ϕ_2 on $\hat{\mathcal{A}}$ supported at two different generators s_1, s_2 with $\phi_1(s_1) = s_1 + v_1$, $\phi_2(s_2) = s_2 + v_2$, and $v_1, v_2 = O(p^2)$. Then the composition $\phi_2 \circ \phi_1$ can also be written as $\phi'_1 \circ \phi'_2$, where ϕ'_1, ϕ'_2 are elementary automorphisms supported at s_1, s_2 with

$$\begin{aligned} \phi'_1(s_1) &= s_1 + \phi_2(v_1) \\ \phi'_2(s_2) &= s_2 + (\phi'_1)^{-1}(v_2); \end{aligned}$$

that is, we can rewrite a composition of elementary automorphisms supported at s_1 and s_2 as a similar composition with the roles of s_1 and s_2 reversed. In addition, if $\phi_1 \equiv \text{Id} \pmod{p^k}$ for some k , then $\phi'_1 \equiv \text{Id} \pmod{p^k}$ and $\phi'_2 \equiv \phi_2 \pmod{p^k}$.

Through this trick, we can rewrite a “partial convergent”

$$(\phi_{-n,2} \cdots \phi_{n,2}) \cdots (\phi_{-n,k} \cdots \phi_{n,k})$$

of the infinite composition as $\phi_{(-n,k)} \cdots \phi_{(n,k)}$, where $\phi_{(j,k)}$ is an elementary automorphism supported at q_j with $\phi_{(j,k)}(q_j) = q_j + v_{j,k}$ for some $v_{j,k} = O(p^2)$. (To this end, note that the composition of two elementary automorphisms supported at q_j is another.) It is easy to see that $v_{j,k} \equiv v_{j,k+1} \pmod{p^k}$ for all k , and thus that $v_{j,k}$ has a limit in $\hat{\mathcal{A}}$ as $k \rightarrow \infty$. Defining $\phi_{(j)}$ for $-n \leq j \leq n$ to be the elementary automorphism supported at q_j with $\phi_{(j)}(q_j) = q_j + \lim_{k \rightarrow \infty} v_{j,k}$ then completes the proof of the lemma. \square

APPENDIX A: ORIENTATION SIGNS

To define the Hamiltonian h and the SFT differential d_{SFT} over \mathbb{Z} , we chose particular orientation signs as shown in Figure 2.4; see also Remark 2.31. These are not the only possible orientation signs leading to a viable LSFT algebra. Here we find all possible combinatorial choices for orientation signs and show that they are all equivalent under basis change. As a corollary, we obtain a refinement of a result in [EES05b]. There two sign choices for Legendrian contact homology in \mathbb{R}^3 are given, one of which recovers the signs from [ENS02] and one of which appears to be different; we show that the two choices are in fact equivalent.

The most general set of orientation signs has eight degrees of freedom $a, b, c, d, e, f, g, h \in \{\pm 1\}$, one for each quadrant of a positive and a negative crossing; see Figure A.1. Note that in the formulation from Section 3, the orientation signs only figure in the definition of h .

To give rise to an LSFT algebra structure, signs must be chosen so that an identity like $\frac{1}{2}\{h, h\} + \delta h = 0$ (Proposition 3.13) holds. In particular, the two terms in $\{h, h\}$ arising from an obtuse disk must cancel. From the proof of Proposition 3.13, we see that we must have $\epsilon\{s', s'^*\}\{s'', s''^*\} = 1$ for all configurations of the form depicted in Figure A.2, where ϵ is the product of the orientation signs over the six shaded corners. One readily deduces that we must have $ab = gh = -ad = -eh$ and $cd = ef = -bc = -fg$, whence $d = -b, g = -e, f = -bc/e$, and $h = ab/e$. This reduces us to four degrees of freedom a, b, c, e .

We can get rid of three further degrees of freedom as follows. Replacing (a, b, c, e) by $(-a, -b, -c, e)$ has the effect in the LSFT algebra of replacing

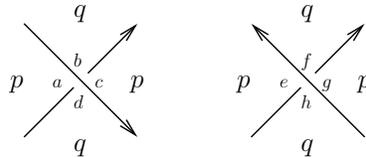


FIGURE A.1. Possible orientation signs for corners. Each of a, b, c, d, e, f, g, h is ± 1 . The left figure is a positive crossing in the usual knot-theoretic sense, with q even and p odd; the right figure is a negative crossing, with q odd and p even.

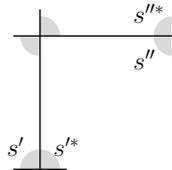


FIGURE A.2. Signs from an obtuse disk (cf. Figure 3.14).

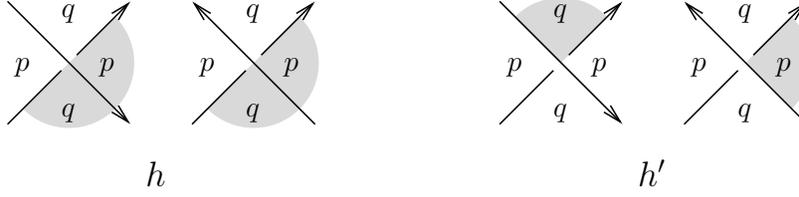


FIGURE A.3. The orientation signs used in this paper (left) and another choice agreeing with [ENS02] (right), resulting in Hamiltonians h and h' .

all p_j, q_j corresponding to positive crossings by $-p_j, -q_j$, thus only modifying the LSFT algebra by a basis change. Similarly, replacing (a, b, c, e) by $(a, b, c, -e)$ just replaces all p_j, q_j corresponding to negative crossings by $-p_j, -q_j$. Furthermore, replacing (a, b, c, e) by $(-a, b, -c, -e)$ has no effect on h or the LSFT algebra, since this simply changes each term w in h by $(-1)^{o(w)}$, where $o(w)$ is the number of odd-degree generators in w and is always even since $|h| = -2$.

Eliminating these three degrees of freedom, we are left with two possibly different equivalence classes of orientation signs, represented by $(a, b, c, e) = (1, 1, -1, 1)$ and $(1, -1, 1, 1)$ and depicted in Figure A.3. These orientation signs yield two Hamiltonians $h, h' \in \hat{\mathcal{A}}$ that agree mod 2. The first is the Hamiltonian used in this paper, satisfying the quantum master equation $\frac{1}{2}\{h, h\} + \delta h = 0$. It can readily be shown that the second satisfies the equation $\frac{1}{2}\{h', h'\} - \delta h' = 0$. We can then define a derivation $d' = \{h, \cdot\} - \delta x$, and $(\hat{\mathcal{A}}, d')$ is an LSFT algebra.

Each of $(\hat{\mathcal{A}}, d)$ and $(\hat{\mathcal{A}}, d')$ induces a choice of signs for the differential on Legendrian contact homology $\mathcal{F}^0 \hat{\mathcal{A}} / \mathcal{F}^1 \hat{\mathcal{A}}$. In [EES05b, Theorem 4.32], two sign rules for Legendrian contact homology are given, essentially corresponding to the two orientations on \mathbb{C} ; these correspond to our $(a, b, c, e) = (1, -1, 1, 1)$ and $(-1, -1, 1, 1)$ and hence to $(\hat{\mathcal{A}}, d')$ and $(\hat{\mathcal{A}}, d)$, respectively. The first sign rule in [EES05b] also agrees with the signs given in [ENS02].⁵

At the time of the writing of [EES05b], it was not known whether the two sign rules led to different contact homology differential graded algebras. In fact, we shall see that they are equivalent. This follows from the corresponding result for the LSFT algebras.

Proposition A.1. *Let Λ be a Legendrian knot in standard contact \mathbb{R}^3 . The LSFT algebra $(\hat{\mathcal{A}}, d)$ for Λ is related by a basis change to the LSFT algebra obtained from $(\hat{\mathcal{A}}, d')$ by conjugation with the involution $t^{\pm 1} \mapsto (-1)^{r(\Lambda)} t^{\pm 1}$, where $r(\Lambda)$ is the rotation number of Λ .*

⁵To translate between our signs and the signs for contact homology in [EES05b, ENS02], we must incorporate the sign $\epsilon(f; s_1)$ (cf. Section 2.3) measuring the orientation of the disk after the p puncture. This has the effect of negating the signs for the corners marked a and g in Figure A.1.

Corollary A.2. *The DGAs given by the two sign rules in [EES05b, Theorem 4.32] are tamely isomorphic if we first replace t in one of the DGAs by $(-1)^{r(\Lambda)}t$. Here the tame isomorphism can be chosen to extend the identity map on the base ring $\mathbb{Z}[t, t^{-1}]$.*

Proof of Proposition A.1. Define an involution ϕ_1 on $\hat{\mathcal{A}}$ that negates p_j, q_j for all j such that p_j, q_j corresponds to a positive crossing, and let $h'' = \phi_1(h')$. The orientation signs defining h'' are the same as those defining h' , except that all signs for positive crossings are reversed. An examination of Figure A.3 then shows that the orientation signs between h and h'' only differ at corners where the knot is oriented into the crossing on both sides of the corner.

Define another involution ϕ_2 on $\hat{\mathcal{A}}$ as follows:

$$\phi_2(s) = \begin{cases} s & \text{if } |s| \equiv 0, 1 \pmod{4} \\ -s & \text{if } |s| \equiv 2, 3 \pmod{4} \\ (-1)^{r(\Lambda)}t^{\pm 1} & \text{if } s = t^{\pm 1}. \end{cases}$$

(Note that the third line is superfluous but has been included for clarity.) We claim that $h'' = -\phi_2(h)$. Indeed, the difference in signs between the appearances of a word w in h and in h'' is $(-1)^{o(w)/2}$, where $o(w)$ is the number of odd s 's appearing in w , so that $o(w)/2$ is the number of corners where the sign changes between h and h'' , cf. Lemma 2.17. Suppose that the word w contains m_j generators (counting multiplicity) whose degree is $j \pmod{4}$ for $j = 0, 1, 2, 3$. Since $|w| = -2$, we have $m_1 + 2m_2 + 3m_3 \equiv 2 \pmod{4}$ and hence

$$\frac{o(w)}{2} = \frac{m_1 + m_3}{2} \equiv m_2 + m_3 + 1 \pmod{2}.$$

The claim follows.

We conclude that $h' = -\phi(h)$ where $\phi = \phi_1 \circ \phi_2$. Note that ϕ is a basis change composed with the map $t^{\pm 1} \mapsto (-1)^{r(\Lambda)}t^{\pm 1}$. By construction, ϕ negates exactly one of each p_j, q_j pair; it follows that $\phi\{x, y\} = -\{\phi(x), \phi(y)\}$ and $\phi(\delta(x)) = -\delta(\phi(x))$ for all x, y . Hence

$$\phi d\phi^{-1}(x) = \phi\{h, \phi^{-1}(x)\} + \phi\delta\phi^{-1}(x) = -\{\phi(h), x\} - \delta(x) = d'(x)$$

and this establishes the proposition. \square

APPENDIX B: STABILIZED KNOTS

In this section, we show that the LSFT algebra of any stabilized Legendrian knot Λ is equivalent via a basis change to an LSFT algebra that

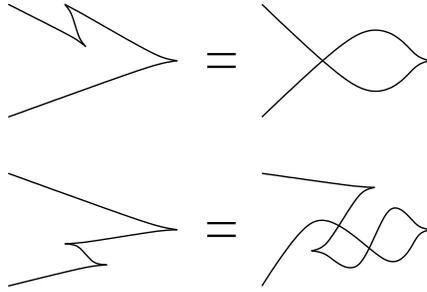


FIGURE B.1. Obtaining a bubble from a stabilization, in the front projection. Which of the two pictures (top or bottom) applies depends on the sign of the stabilization and the orientation of the rightmost cusp of the front.

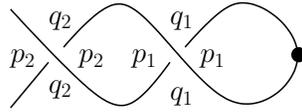


FIGURE B.2. The xy projection of a stabilized knot Λ .

depends only on $tb(\Lambda)$ and $r(\Lambda)$, the classical Legendrian invariants associated to Λ . This implies that the LSFT algebra of a stabilized knot contains no interesting information about the knot.

Let Λ be a stabilized knot, i.e., a knot Legendrian isotopic to one whose front diagram contains a zigzag. Up to isotopy, we can assume that the front of Λ has a zigzag next to its rightmost cusp. We can further isotop the front to obtain a “bubble” at the rightmost cusp; see Figure B.1. The resolution of a bubble is shown in Figure B.2. It follows that, up to equivalence of LSFT algebras, we can assume that the xy diagram for Λ , given by resolving its front, contains the piece shown in Figure B.2, and no part of the diagram lies further to the right than the depicted part.

With q_1, p_1, q_2, p_2 as labeled and the base point \bullet as shown in Figure B.2, the LSFT algebra for Λ satisfies

$$\begin{aligned} F_d &= p_1 \\ d(p_1) &= 0 \\ d(q_1) &= 1 - p_2 + q_1 p_1 q_1. \end{aligned}$$

(These signs are correct if the rightmost loop in Figure B.2 is oriented counterclockwise; if it is oriented clockwise, then $d(q_1) = -1 + p_2 - q_1 p_1 q_1$, but this is equivalent to the given signs after we replace q_1 by $-q_1$.) By further conjugating by the basis change given by the elementary automorphism $p_2 \mapsto -p_2 + q_1 p_1 q_1$, we can write $d(q_1) = 1 + p_2$ instead.

Definition B.1. The LSFT algebra $(\hat{\mathcal{A}}, d)$ generated by $q_1, \dots, q_n, p_1, \dots, p_n, t^{\pm 1}$ is of *ordered stabilized type* if it has the form:

$$\begin{aligned} F_d &= p_1 \\ d(p_1) &= 0 \\ d(q_1) &= 1 + p_2 \\ d(p_2) &= [p_1, q_1] \\ d(q_j) &= (\text{function of } t^{\pm 1}, q_1, \dots, q_{j-1}) + O(p) \\ d(p_j) &= (\text{function of } t^{\pm 1}, q_1, \dots, q_n, p_1, p_2, p_{j+1}, \dots, p_n) + O(p^2) \\ d(t) &= O(p) \end{aligned}$$

where $j = 1, \dots, n$.

Lemma B.2. *Up to equivalence, the LSFT algebra for a stabilized knot is of ordered stabilized type.*

Proof. From the discussion prior to Definition B.1, we can write the LSFT algebra of any stabilized knot to satisfy $F_d = p_1$, $d(p_1) = 0$, $d(q_1) = 1 + p_2$, and hence $d(p_2) = [p_1, q_1]$ as well. Now order the crossings of the xy projection to satisfy $h(q_3) \leq h(q_4) \leq \dots \leq h(q_n)$, where h is the height of the corresponding Reeb chord. The remainder of the conditions in Definition B.1 then follow automatically from Lemma 2.22. \square

A linear map $f: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ is a map that sends $y \in \hat{\mathcal{A}}$ to a sum of terms, each of which includes y once. Given a linear map f and $y \in \hat{\mathcal{A}}$, we can define another linear map $D_y f: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ as the ‘‘derivative’’ of f : $(D_y f)(z) = d(f(y))|_{d(y)=z}$. Thus if $f(y) = x_1 y x_2$ for $x_1, x_2 \in \hat{\mathcal{A}}$, then

$$(D_y f)(z) = d(x_1) y x_2 + (-1)^{|x_1|} x_1 z x_2 + (-1)^{|x_1 y|} x_1 y d(x_2).$$

Let $(\hat{\mathcal{A}}, d)$ be of ordered stabilized type. For $k \geq -1$, define a linear map $f_k: \hat{\mathcal{A}} \rightarrow \hat{\mathcal{A}}$ inductively as follows: $f_{-1}(y) = f_0(y) = 0$ and

$$f_k(y) = f_{k-1}(y) + q_1 [p_1, y] - q_1 (D_y f_{k-1})(f_{k-1}(y)) \pmod{p^{k+1}},$$

where $\pmod{p^{k+1}}$ indicates that we drop all terms of order $k+1$ or higher in the p 's (i.e., in $\mathcal{F}^{k+1} \hat{\mathcal{A}}$ for all y). The first few f_k 's are given as follows:

$$\begin{aligned} f_0(y) &= 0 \\ f_1(y) &= q_1 [p_1, y] \\ f_2(y) &= q_1 [p_1, y] + q_1^2 [p_1, q_1 [p_1, y]] - q_1 p_2 [p_1, y] \\ f_3(y) &= q_1 [p_1, y] + q_1^2 [p_1, q_1 [p_1, y]] - q_1 p_2 [p_1, y] + q_1^2 [p_1, q_1^2 [p_1, q_1 [p_1, y]]] \\ &\quad - q_1^3 [p_1, p_2 [p_1, y]] - q_1^2 [p_1, q_1 p_2 [p_1, y]] - q_1 p_2 q_1 [p_1, q_1 [p_1, y]] + q_1 p_2^2 [p_1, y]. \end{aligned}$$

Lemma B.3. *For all k , $f_k(y)$ is a linear function in y with coefficients involving only p_1, q_1, p_2 , and $f_k(y) \in \mathcal{F}^1 \hat{\mathcal{A}}$ (i.e., every term in $f_k(y)$ involves some p). Furthermore, $f_k(y)$ does not depend on $(\hat{\mathcal{A}}, d)$.*

Proof. Clear by induction. \square

Now suppose that $(\hat{\mathcal{A}}, d)$ is of ordered stabilized type. Let s be a generator of $\hat{\mathcal{A}}$ besides p_1, q_1, p_2 . We inductively define a sequence of differentials d_k , elements $x_k \in \hat{\mathcal{A}}$, and elementary automorphisms ϕ_k supported at s , as follows:

- $d_{-1} = d$;
- $x_k = d_{k-1}(s) - f_{k-1}(s) \pmod{p^{k+1}}$;
- $\phi_k(s) = s - q_1 x_k$;
- $d_k = \phi_k^{-1} \circ d_{k-1} \circ \phi_k$.

Note that it is not clear a priori that ϕ_k is an elementary automorphism (in particular, invertible).

Lemma B.4. *For all $k \geq 0$, we have:*

- $d_k(x) = d(x)$ for $x = p_1, q_1, p_2$ and $F_{d_k} = p_1$;
- $x_k = O(p^k)$;
- ϕ_k is an elementary automorphism;
- $d_k(s) = f_k(s) + O(p^{k+1})$.

Proof. We prove the lemma by induction. For $k = 0$, since d is ordered, $x_0 = d(s) \pmod{p}$ does not involve s , so ϕ_0 is elementary. We then have $d(x_0) = d^2(s) + O(p) = [p_1, s] + O(p) = O(p)$ and thus

$$d_0(s) = \phi_0^{-1}(d(s - q_1 x_0)) = \phi_0^{-1}(x_0 - x_0 + q_1 d(x_0) + O(p)) = \phi_0^{-1}(O(p)) = O(p).$$

Now assume the lemma holds for $k-1 \geq 0$. Since ϕ_k is supported at s and $d_{k-1}(x)$ does not involve s for $x = p_1, q_1, p_2$, it follows that $d_k(x) = d_{k-1}(x)$ for these values of x , while $F_{d_k} = \phi_k^{-1}(F_{d_{k-1}}) = p_1$.

Next, since $d_{k-1}(s) = f_{k-1}(s) + O(p^k)$, $x_k = O(p^k)$ by definition. It follows that ϕ_k is elementary. (If s is a p and $k = 1$, then $x_0 = 0$ and $x_1 = d(s) \pmod{p^2}$; since d is ordered, x_1 does not involve s and so ϕ_k is elementary.)

Finally we check that $d_k(s) = f_k(s) + O(p^{k+1})$. By Lemma B.3, $f_{k-1}(s)$ is linear in s with coefficients involving only p_1, q_1, p_2 ; since $d_{k-1} = d$ on p_1, q_1, p_2 , we have $d_{k-1}(f_{k-1}(s)) = (D_s f_{k-1})(d_{k-1}(s))$. By the induction assumption, $d_{k-1}(s) \equiv f_{k-1}(s) \pmod{p^k}$; since every term in f_{k-1} involves a p by Lemma B.3 again, we conclude that

$$d_{k-1}(f_{k-1}(s)) = (D_s f_{k-1})(d_{k-1}(s)) \equiv (D_s f_{k-1})(f_{k-1}(s)) \pmod{p^{k+1}}.$$

Thus

$$\begin{aligned} d_{k-1}\phi_k(s) &= d_{k-1}(s - q_1 x_k) = d_{k-1}(s) - (1 + p_2)x_k + q_1 d_{k-1}(x_k) \\ &= f_{k-1}(s) + q_1(d_{k-1}^2(s) - d_{k-1}f_{k-1}(s)) + O(p^{k+1}) \\ &= f_{k-1}(s) + q_1[p_1, s] - q_1(D_s f_{k-1})(f_{k-1}(s)) + O(p^{k+1}) \\ &= f_k(s) + O(p^{k+1}). \end{aligned}$$

Now ϕ_k is the identity mod p^k and $f_k(s) \in \mathcal{F}^1 \hat{\mathcal{A}}$, whence $d_k(s) = \phi_k^{-1} d_{k-1} \phi_k(s) = \phi_k^{-1}(f_k(s)) + O(p^{k+1}) = f_k(s) + O(p^{k+1})$. \square

Lemma B.5. *For all k , we have $f_k(s) = f_{k-1}(s) + O(p^k)$.*

Proof. This follows directly from Lemma B.4. Since ϕ_k is the identity mod p^k , $d_k(s) = d_{k-1}(s) + O(p^k)$. From Lemma B.4 again, $f_k(s) = d_k(s) + O(p^{k+1})$ and $f_{k-1}(s) = d_{k-1}(s) + O(p^k)$. The lemma follows. \square

Because of Lemma B.5, we can define an element $f(s) \in \hat{\mathcal{A}}$ to be the limit $\lim_{k \rightarrow \infty} f_k(s)$.

Lemma B.6. *If $(\hat{\mathcal{A}}, d)$ is of ordered stabilized type and s is one of $q_2, \dots, q_n, p_3, \dots, p_n, t$, then there is an elementary automorphism ϕ supported at s such that if $d' = \phi^{-1} \circ d \circ \phi$, then $(\hat{\mathcal{A}}, d')$ is of ordered stabilized type and $d'(s) = f(s)$.*

Proof. Using Lemma B.4, set $\phi = \lim_{k \rightarrow \infty} (\phi_0 \circ \phi_1 \circ \dots \circ \phi_k)$, a well-defined limit since ϕ_k is the identity mod p^k for all k . \square

Proposition B.7. *An LSFT algebra of ordered stabilized type is equivalent under basis change to the LSFT algebra whose differential is given by*

$$\begin{aligned} F_d &= p_1 \\ d(p_1) &= 0 \\ d(q_1) &= 1 + p_2 \\ d(p_2) &= [p_1, q_1] \\ d(q_j) &= f(q_j) \text{ for } 2 \leq j \leq n \\ d(p_j) &= f(p_j) \text{ for } 3 \leq j \leq n \\ d(t) &= f(t). \end{aligned}$$

Proof. Successively apply Lemma B.6 with $s = q_2, \dots, q_n, p_3, \dots, p_n, t$. \square

For any Legendrian knot Λ , the LSFT algebra of Λ encodes both the Thurston–Bennequin and rotation numbers of Λ : $r(\Lambda)$ is $-1/2$ times the degree of t , while $tb(\Lambda)$ is the difference between the number of q generators of even degree and the number of odd degree. However, if Λ is stabilized, Proposition B.7 implies that the LSFT algebra encodes nothing else.

Corollary B.8. *The LSFT algebra of a stabilized knot Λ is equivalent to an LSFT algebra depending only on $tb(\Lambda)$ and $r(\Lambda)$.*

REFERENCES

- [BEE] Frédéric Bourgeois, Tobias Ekholm, and Yakov Eliashberg. Effect of Legendrian surgery. arXiv:0911.0026.
- [Che02] Yuri Chekanov. Differential algebra of Legendrian links. *Invent. Math.*, 150(3):441–483, 2002.

- [CL] Kai Cieliebak and Janko Latschev. The role of string topology in symplectic field theory. arXiv:0706.3284.
- [CL06] Octav Cornea and François Lalonde. Cluster homology: an overview of the construction and results. *Electron. Res. Announc. Amer. Math. Soc.*, 12:1–12 (electronic), 2006.
- [CS] Moira Chas and Dennis Sullivan. String topology. math.GT/9911159.
- [EES05a] Tobias Ekholm, John Etnyre, and Michael Sullivan. The contact homology of Legendrian submanifolds in \mathbb{R}^{2n+1} . *J. Differential Geom.*, 71(2):177–305, 2005. math.SG/0210124.
- [EES05b] Tobias Ekholm, John Etnyre, and Michael Sullivan. Orientations in Legendrian contact homology and exact Lagrangian immersions. *Internat. J. Math.*, 16(5):453–532, 2005. math.SG/0408411.
- [EES07] Tobias Ekholm, John Etnyre, and Michael Sullivan. Legendrian contact homology in $P \times \mathbb{R}$. *Trans. Amer. Math. Soc.*, 359(7):3301–3335 (electronic), 2007. math.SG/0505451.
- [EGH00] Y. Eliashberg, A. Givental, and H. Hofer. Introduction to symplectic field theory. *Geom. Funct. Anal.*, (Special Volume, Part II):560–673, 2000. GAFA 2000 (Tel Aviv, 1999).
- [Eli98] Yakov Eliashberg. Invariants in contact topology. In *Proceedings of the International Congress of Mathematicians, Vol. II (Berlin, 1998)*, number Extra Vol. II, pages 327–338 (electronic), 1998.
- [Eli07] Yakov Eliashberg. Symplectic field theory and its applications. In *International Congress of Mathematicians. Vol. I*, pages 217–246. Eur. Math. Soc., Zürich, 2007.
- [ENS02] John B. Etnyre, Lenhard L. Ng, and Joshua M. Sabloff. Invariants of Legendrian knots and coherent orientations. *J. Symplectic Geom.*, 1(2):321–367, 2002. math.SG/0101145.
- [Etn05] John B. Etnyre. Legendrian and transversal knots. In *Handbook of knot theory*, pages 105–185. Elsevier B. V., Amsterdam, 2005. math.SG/0306256.
- [Fer02] Emmanuel Ferrand. On Legendrian knots and polynomial invariants. *Proc. Amer. Math. Soc.*, 130(4):1169–1176 (electronic), 2002. math.GT/0002250.
- [Ng01] Lenhard L. Ng. Invariants of Legendrian links. Ph.D. dissertation, MIT, 2001.
- [Ng03] Lenhard L. Ng. Computable Legendrian invariants. *Topology*, 42(1):55–82, 2003. math.GT/0002250.
- [Ng08] Lenhard Ng. Framed knot contact homology. *Duke Math. J.*, 141(2):365–406, 2008. math/0407071.
- [NT04] Lenhard Ng and Lisa Traynor. Legendrian solid-torus links. *J. Symplectic Geom.*, 2(3):411–443, 2004. math.SG/0407068.
- [Pos93] L. E. Positsel’skiĭ. Nonhomogeneous quadratic duality and curvature. *Funktsional. Anal. i Prilozhen.*, 27(3):57–66, 96, 1993.
- [Sab03] Joshua M. Sabloff. Invariants of Legendrian knots in circle bundles. *Commun. Contemp. Math.*, 5(4):569–627, 2003. math.SG/0208214.
- [Sch99] Albert Schwarz. Noncommutative supergeometry and duality. *Lett. Math. Phys.*, 50(4):309–321, 1999.

MATHEMATICS DEPARTMENT, DUKE UNIVERSITY, DURHAM, NC 27708

E-mail address: ng@math.duke.edu